

# High-Dimensional Stochastic Optimal Control using Continuous Tensor Decompositions

Alex Gorodetsky

Sertac Karaman

Youssef Marzouk

**Abstract**—Motion planning and control problems are embedded and essential in almost all robotics applications. These problems are often formulated as stochastic optimal control problems and solved using dynamic programming algorithms. Unfortunately, most existing algorithms that guarantee convergence to optimal solutions suffer from the *curse of dimensionality*: the run time of the algorithm grows exponentially with the dimension of the state space of the system. We propose novel dynamic programming algorithms that alleviate the curse of dimensionality in problems that exhibit certain *low-rank* structure. The proposed algorithms are based on continuous tensor decompositions [24] recently developed by the authors. Essentially, the algorithms represent high-dimensional functions (e.g., the value function) in a compressed format, and directly perform dynamic programming computations (e.g., value iteration, policy iteration) in this format. Under certain technical assumptions, the new algorithms guarantee convergence towards optimal solutions with arbitrary precision. Furthermore, the run times of the new algorithms scale polynomially with the state dimension and polynomially with the ranks of the value function. This approach realizes substantial computational savings in “compressible” problem instances, where value functions admit low-rank approximations. We demonstrate the new algorithms in a wide range of problems, including a simulated six-dimensional agile quadcopter maneuvering example and a seven-dimensional aircraft perching example. In some of these examples, we estimate computational savings of up to ten orders of magnitude over standard value iteration algorithms. We further demonstrate the algorithms running in real time on board a quadcopter during a flight experiment under motion capture.

**Index Terms**—stochastic optimal control, motion planning, dynamic programming, tensor decompositions

## I. INTRODUCTION

The control synthesis problem is to find a feedback control law, or controller, that maps each state of a given dynamical system to its control inputs, often optimizing given performance or robustness criteria [5], [41]. Control synthesis problems are prevalent in several robotics applications, such as agile maneuvering [44], humanoid robot motion control [17], [18], and robot manipulation [56], just to name a few.

Analytical approaches to control synthesis problems make simplifying assumptions on the problem setup to derive explicit formulas that determine controller parameters. Common assumptions include dynamics described by linear ordinary differential equations and Gaussian noise. In most cases, these assumptions are so severe that analytical approaches find little direct use in robotics applications.

On the other hand, computational methods for control synthesis can be formulated for a fairly large class of dy-

namical systems [4], [5], [52]. However, unfortunately, most control synthesis problems turn out to be prohibitively computationally challenging, particularly for systems with high-dimensional state spaces. In fact, Bellman [2] coined the term *curse of dimensionality* in 1961 to describe the fact that the computational requirements grow exponentially with increasing dimensionality of the state space of the system.

In this paper, we propose a novel class of computational methods for stochastic optimal control problems. The new algorithms are enabled by a novel representation of the controller that allows efficient computation of the controller. This new representation can be viewed as a type of “compression” of the controller. The compression is enabled by a continuous tensor decomposition method, called the *function train*, which was recently proposed by the authors [24] as a continuous analogue of the well-known tensor-train decomposition [47], [49]. Our algorithms result in control synthesis problems with run time that scales polynomially with the dimension and the rank of the optimal value function. These control synthesis algorithms run several orders of magnitude faster than standard dynamic programming algorithms, such as value iteration. The resulting controllers also require several orders of magnitude less storage.

### A. Related work

Computational hurdles are present in most decision making problems in the robotics domain. A closely related problem is motion planning: the problem of finding a dynamically-feasible, collision-free trajectory from an initial configuration to a final configuration for a robot operating in a complex environment. Motion planning problems are embedded and essential in almost all robotics applications, and they have received significant attention since the early days of robotics research [40], [41]. However, it is well known that these problems are computationally challenging [10]. For instance, a simple version of the motion planning problem is PSPACE-hard [10]. In other words, it is unlikely that there exists a complete algorithm with running time that scales polynomially with increasing degrees of freedom, i.e., the dimensionality of the configuration space of the robot. In fact, the run times of all known complete algorithms scale exponentially with dimensionality [10], [41]. Most of these algorithms construct a discrete abstraction of the continuous configuration space, the size of which scales exponentially with dimensionality.

Yet, there are several practical algorithms for motion planning, some of which even provide completeness properties. For instance, a class of algorithms called sampling-based

The authors are with the Department of Aeronautics and Astronautics, Massachusetts Institute of Technology.

algorithms [30], [34], [41], [42] construct a discrete abstraction, often called a roadmap, by sampling the configuration space and connecting the samples with dynamically-feasible, collision-free trajectories. The result is a class of algorithms that find a feasible solution, when one exists, in a reasonable amount of time for many problem instances, particularly for those that have good “visibility” properties [31], [33]. These algorithms provide probabilistic completeness guarantees, i.e., they return a solution, when one exists, with probability approaching one as the number of samples increases.

In the same way, most practical approaches to motion planning avoid the construction of a grid to prevent intractability. Instead, they construct a “compact” representation of the continuous configuration space. The resulting compact data structure not only provides substantial computational gains, but also it still accurately represents the configuration space in a large class of problem instances, e.g., those with good visibility properties. These claims can be made precise in provable guarantees such as probabilistic completeness and the exponential of rate of decay of the probability of failure.

It is worth noting at this point that optimal motion planning, i.e., the problem of finding a dynamically-feasible, collision-free trajectory that minimizes some cost metric, has also been studied widely [32]. In particular, sampling-based algorithms have been extended to optimal motion planning problems recently [32]. Various trajectory optimization methods have also been developed and demonstrated [53], [63]. Another relevant problem that attracted attention recently is feedback motion planning, in which the goal is to synthesize a feedback control often by generating controllers that track motion [57].

The algorithm proposed in this paper is a different, novel approach to stochastic optimal control problems, which provides significant computational savings with provable guarantees.

## B. Tensor decomposition methods

In this paper, we propose a new class of algorithms for high-dimensional instances of stochastic optimal control problems that are based on compressing associated value functions. Moreover, we compress a *functional*, rather than a discretized, representation of the value function. This approach enables fast evaluation of the value function at arbitrary points in the state space, without any decompression. Our approach offers orders of magnitude reduction in the required storage costs.

Specifically, the new algorithms are based on the function train (FT) decomposition [24] recently proposed by the authors. Multivariate functions in the FT format are represented by a set of matrix-valued functions whose sizes correspond to the FT rank. Hence, low-rank multivariate functions can be represented in the FT format with a few parameters, and in this paper we use this compression to represent the value function of an associated stochastic optimal control problem.

In addition to *representing* the value function, we *compute* with value functions directly in compressed form; no decompression is ever performed. Specifically, we create compressed versions of value iteration, policy iteration, and other algorithms for solving dynamic programming problems. These DP problems are obtained through consistent discretizations of

continuous-time continuous-space stochastic optimal control problems obtained using the Markov chain approximation (MCA) method [39]. Note that even though the MCA method relies on discretization, the FT still allows us to maintain a *functional* representation, valid for any state within the state space.

As a result, the new algorithms provide substantial computational gains in terms of both computation time and storage space, when compared to standard dynamic programming methods.

The proposed algorithms exploit the *low-rank* structure commonly found in *separable* functions. This type of structure has been widely exploited within numerical analysis literature on tensor decompositions [25], [26], [35]. Indeed the FT decomposition itself is an extension of tensor train (TT) decomposition developed by Oseledets [47], [49]. The TT decomposition works with discrete *arrays*; it represents a  $d$ -dimensional array as a multiplication of  $d$  matrices. The FT decomposition, on the other hand, works directly with *functions*. This representation allows a wider variety of possible operations to be performed in FT format, e.g., integration, differentiation, addition, and multiplication of low-rank functions. These operations are problematic in the purely discrete framework of tensor decompositions since element-wise operation with tensors is undefined, e.g., one cannot add arrays with different number of elements.

We use the term *compressed continuous computation* for such FT-based numerical methods [24]. Compressed continuous computation algorithms can be considered an extension of the *continuous computation* framework to high-dimensional function spaces via the FT-based compressed representation. The continuous computation framework, roughly referring to computing with directly with functions (as opposed to discrete arrays), was first realized by Chebfun, a Matlab software package developed by Trefethen, Battles, Townsend, Platte, and others [50]. In this software package, the user computes with univariate [1], [50], bivariate [58], and trivariate [27] functions that are represented in Chebyshev polynomial bases. Our recent work [23] extended this framework to the general multivariate case by building a bridge between continuous computation and low-rank tensor decompositions. Our prior work has resulted in the software package, called Compressed Continuous Computation ( $C^3$ ) [21], implemented in the C programming language. Examples presented in this paper use the  $C^3$  software package, and they are available online on GitHub [22].

In short, our compressed continuous computation framework leverages both the advantages of continuous computation and low-rank tensor decompositions. The advantage of *compression* is tractability when working directly with high-dimensional computational structures. For instance, a seven-dimensional array with one hundred points in each dimension includes trillion points in total. As a result, even the storage space required cannot be satisfied by any existing computer, let alone the computation times. The computational requirements increase rapidly with increasing dimensionality.

The advantages of the functional, or *continuous*, representation (over the array-based TT) include the ability to compare

value functions resulting from different discretization levels and the ability to evaluate these functions outside of some discrete set of nodes. We leverage these advantage in this paper in two ways: (i) we develop multi-level schemes based on low-rank prolongation and interpolation operators; (ii) we evaluate optimal policies for any state in the state space during the execution of the controller.

### C. Contributions

The main contributions of this paper are as follows. First, we propose novel compressed continuous computation algorithms for dynamic programming. Specifically, we utilize the function train decomposition algorithms to design FT-based value iteration, policy iteration, and one-way multigrid algorithms. These algorithms work with a Markov chain approximation for a given continuous-time continuous-space stochastic optimal control problem. They utilize the FT-based representation of the value function to map the discretization due to Markov chain approximation into a compressed, functional representation.

Second, we prove that, under certain conditions, the new algorithms guarantee convergence to optimal solutions, whenever the standard dynamic programming algorithms also guarantee convergence for the same problem instance. We also prove upper bounds on computational requirements. In particular, we show that the run time of the new algorithms scale polynomially with dimension and polynomially with the rank of the value function, while even the storage requirements for existing dynamic programming algorithms clearly scale exponentially with dimension.

Third, we demonstrate the new algorithms in challenging problem instances. In particular, we consider perching problem that features non-linear non-holonomic non-control-affine dynamics. We estimate that the computational savings reach roughly ten orders of magnitude. In particular, the controller that we find fits in roughly 1MB of space in the compressed FT format; we estimate that a full look up table, for instance, one computed using standard dynamic programming algorithms, would have required around 20 TB of memory.

We also consider the problem of maneuvering a quadcopter through a small window. This leads to a six-dimensional non-linear non-holonomic non-control-affine stochastic optimal control problem. We compute the near-optimal solution. We demonstrate the resulting controller in both simulation and experiment. In experiment, we utilize a motion capture system for full state information, and run the resulting controller in real time on board the vehicle.

### D. Organization

The paper is organized as follows. We introduce the stochastic optimal control problem in Section II and the Markov chain approximation method in Section III. We briefly describe the compressed continuous computation framework in Section III. We describe the proposed algorithms in Section V. We analyze their convergence properties and their computational costs in Section VI. We discuss a wide range of numerical examples and experiments in Section VII. Section VIII offers some concluding remarks.

## II. STOCHASTIC OPTIMAL CONTROL

In this section, we formulate a class of continuous-time continuous-space stochastic optimal control problems. Background is provided in Section II-A. Under some mild technical assumptions, the optimal control is a Markov policy, i.e., a mapping from the state space to the control space, that satisfies the Hamilton-Jacobi-Bellman (HJB) equation. We introduce the notion of Markov policies and the HJB equation in Sections II-B and II-C, respectively.

### A. Stochastic optimal control

Denote the set of integers and the set of reals by  $\mathbb{Z}$  and  $\mathbb{R}$ , respectively. We denote the set of all positive real numbers by  $\mathbb{R}_+$ . Similarly, the set of positive integers is denoted by  $\mathbb{Z}_+$ . Let  $d, d_u, d_w \in \mathbb{Z}_+$ ,  $\mathcal{X} \subset \mathbb{R}^d$  and  $\mathcal{U} \subset \mathbb{R}^{d_u}$  be compact sets with smooth boundaries and non-empty interiors,  $\mathcal{T} \subset \mathbb{R}_+$ , and  $\{w(t) : t \geq 0\}$  be a  $d_w$ -dimensional Brownian motion defined on some probability space  $(\Omega, \mathcal{F}, \mathbb{P})$ , where  $\Omega$  is a sample space,  $\mathcal{F}$  is a  $\sigma$ -algebra, and  $\mathbb{P}$  is a probability measure.

Consider a dynamical system described by the following stochastic differential equation in the differential form:

$$dx(t) = B(t, x(t), u(t))dt + \mathcal{D}(t, x(t))dw(t), \text{ for all } t \in \mathcal{T}, \quad (1)$$

where  $B : \mathcal{T} \times \mathcal{X} \times \mathcal{U} \rightarrow \mathbb{R}^d$  is a vector-valued function, called the *drift*, and  $\mathcal{D} : \mathcal{T} \times \mathcal{X} \rightarrow \mathbb{R}^{d \times d_w}$  is a matrix-valued function, called the *diffusion*. Strictly speaking, for any admissible control process<sup>1</sup>  $\{u(t) : t \geq 0\}$ , the solution to this differential form is a stochastic process  $\{x(t) : t \geq 0\}$  satisfying the following integral equation: For all  $t \in \mathcal{T}$ ,

$$x(t) = x(0) + \int_0^t B(\tau, x(\tau), u(\tau)) d\tau + \int_0^t \mathcal{D}(\tau, x(\tau), u(\tau)) dw(\tau), \quad (2)$$

where the last term on the right hand side is the usual Itô integral [46]. We assume that the drift and diffusion are measurable, continuous, and bounded functions. These conditions guarantee existence and uniqueness of the solution to Equation (2) [46].

We focus on two problem formulations: the fixed-cost finite-horizon problem and the discounted-cost infinite-horizon problem. Our description of these problems and the corresponding notation closely follows that of Fleming and Soner [19].

Let  $\mathcal{O} \subset \mathcal{X}$  denote an open subset. If  $\mathcal{O} \neq \mathbb{R}^d$  then let its boundary  $\partial\mathcal{O}$  be a compact  $(d-1)$ -dimensional manifold of class  $C^3$ , i.e., the set of 3-times differential functions. Let  $g, \psi$  denote continuous stage and terminal cost functions, respectively, that satisfy polynomial growth conditions:

$$|g(t, x, u)| \leq C(1 + |x|^k + |u|^k), \\ |\psi(t, x)| \leq C(1 + |x|^k),$$

<sup>1</sup>Suppose the control process  $\{u(t) : t \geq 0\}$  is defined on the same probability space  $(\Omega, \mathcal{F}, \mathbb{P})$  which the Wiener process  $\{w(t) : t \geq 0\}$  is also defined on. Then,  $\{u(t) : t \geq 0\}$  is said to be *admissible* with respect to  $\{w(t) : t \geq 0\}$ , if there exists a filtration  $\{\mathcal{F}_t : t \geq 0\}$  defined on  $(\Omega, \mathcal{F}, \mathbb{P})$  such that  $u(t)$  is  $\mathcal{F}_t$ -adapted and  $w(t)$  is an  $\mathcal{F}_t$ -Wiener process. Kushner et al. [39] provide the precise measure theoretic definitions.

for some constants  $C \in \mathbb{R}, k \in \mathbb{N}$ .

In the finite-cost finite-horizon stochastic optimal control problem, we seek a control law that acts upon the system until a finite time is reached or until the state exits the space  $\mathcal{O}$ . Define  $\mathcal{T} = [t_0, t_1]$  for fixed  $t_0, t_1 \in \mathbb{R}_+$  with  $t_0 < t_1 < \infty$ . Let  $\mathcal{Q} = [t_0, t_1] \times \mathcal{O}$  denote the product space of times and states. Then, the *exit time*  $\tau$  of the system is defined to be

$$\tau = \inf\{s : (s, x(s)) \notin \mathcal{Q}\}, \quad s > t.$$

For example, we have  $\tau = t_1$ , if for all times  $t \leq s \leq t_1$  the state remains in the interior of the state space, i.e.,  $x(s) \in \mathcal{O}$ . For any product state  $(t, z) \in \mathcal{Q}$  and admissible process  $u = \{u(t) : t \geq 0\}$ , the cost functional is

$$c(t, z; u) = \mathbb{E} \left[ \int_t^\tau g(s, x(s), u(s)) ds + \psi(\tau, x(\tau)) \mid x(t) = z \right].$$

This functional describes the expected cost incurred under control  $u$  for a system initialized with initial condition  $(t, z)$ .

The *finite-cost finite-horizon stochastic optimal control problem* is to find a control  $u(t)$  such that  $c(t, z; u)$  is minimized for all  $z \in \mathcal{O}$ , subject to Equation (1).

In the discounted-cost infinite-horizon stochastic optimal control problem, we consider only time-invariant dynamical systems where Equation (1) is modified according to

$$dx(t) = B(x(t), u(t))dt + D(x(t), u(t))dw(t), \quad x(0) = z,$$

where for  $z \in \mathcal{O}$ . In this case, the exit time must also be modified because the system could evolve indefinitely. Define the *exit time*  $\tau$  as either the first time that the state  $x(s)$  exits from  $\mathcal{O}$ , or we set  $\tau = \infty$  if the state remains forever within  $\mathcal{O}$ , i.e.,  $x(s) \in \mathcal{O}$  for all  $s \geq 0$ . Within this formulation, we can still use a terminal cost  $\psi$  for the cases when  $\tau < \infty$ . To accommodate finite exit times, we use the indicator function  $\chi_{\tau < \infty}$  that evaluates to one if the state exits  $\mathcal{O}$  and to zero otherwise. The cost functional is defined as:

$$\bar{c}(z; u) = \mathbb{E} \left[ \int_0^\tau e^{-\beta s} g(s, x(s), u(s)) ds + \chi_{\tau < \infty} e^{-\beta \tau} \psi(\tau, x(\tau)) \right], \quad x(0) = z,$$

where  $\beta > 0$  is a discount factor.

The *discounted-cost infinite-horizon stochastic optimal control problem* is to find a control  $u(t)$  such that  $\bar{c}(z; u)$  is minimized for all  $z \in \mathcal{O}$ , subject to Equation (1). We require that the cost until exit is bounded

$$\mathbb{E} \left[ \int_0^\tau \exp^{-\beta s} |g(s, x(s), u(s))| ds \right] \leq \infty,$$

for this problem to be well defined [19].

### B. Markovian policies

A *Markov policy* is a mapping  $\mu : \mathcal{T} \times \mathcal{X} \rightarrow \mathcal{U}$  that assigns a control input to each time and state. Under a Markov policy  $\mu$ , an admissible control is obtained according to  $u(t) = \mu(t, x(t))$ . For both problem formulations described in the previous section, we can correspondingly denote the

cost functional associated with a specific Markov policy  $\mu$  to be

$$c_\mu(t, z) = \mathbb{E} \left[ \int_t^\tau g(s, x(s), \mu(s, x(s))) ds + \psi(\tau, x(\tau)) \right], \quad x(t) = z,$$

for the finite-cost finite-horizon problem, and

$$\bar{c}_\mu(z) = \mathbb{E} \left[ \int_0^\tau e^{-\beta s} g(s, x(s), \mu(s, x(s))) ds + \chi_{\tau < \infty} e^{-\beta \tau} \psi(x(\tau)) \right], \quad x(0) = z,$$

for the discounted-cost infinite-horizon problem.

Under certain conditions, one can show that a Markov policy is at least as good as any other arbitrary  $\mathcal{F}_t$ -adapted policy; see for example Theorem 11.2.3 by Øksendal [46]. In this work we assume these conditions hold, and only work with Markov control policies. Storing Markov control policies allows us to avoid storing trajectories of the system when considering what action to apply. Instead, Markov policies only require knowledge of the current time and state and are computationally efficient to use in practice.

Let  $y = (t, z)$  denote the product state of the finite-cost finite-horizon problem, and let  $y = z$  refer to the state of the discounted-cost infinite-horizon problem. Then, the stochastic control problem is to find an *optimal* cost  $\bar{c}_{\mu^*}$  with the following property

$$\bar{c}_{\mu^*}(y) = \inf_{\mu} \bar{c}_\mu(y), \quad \text{for all } y,$$

subject to Equation (1).

### C. Dynamic programming

We can formulate the stochastic optimal control problem as a dynamic programming problem. In the dynamic programming formulation, we seek an optimal value function  $v(t, z)$  defined as

$$v(t, z) = \inf_{\mu} c_\mu(t, z) \text{ for all } z \in \mathcal{O}.$$

For continuous-time continuous-space stochastic optimal control problems, the optimal value function satisfies a partial differential equation, called the Hamilton-Jacobi-Bellman (HJB) partial differential equation (PDE) [19].

To define the HJB PDE, we first introduce some notation. Let  $\mathcal{S}_+^d$  denote the set of symmetric, nonnegative definite matrices. Let  $\mathbf{A} \in \mathcal{S}_+^d$  and  $\mathcal{A} = \mathcal{D}\mathcal{D}^T$ , then the trace  $\text{tr } \mathcal{A}\mathbf{A}$  is defined as

$$\text{tr } \mathcal{A}\mathbf{A} = \sum_{i,j}^d \mathcal{A}[i, j] \mathbf{A}[i, j].$$

For  $(t, z) \in \mathcal{Q}$ ,  $p \in \mathcal{X}$ ,  $\mathbf{A} \in \mathcal{S}_+^d$ , define the *Hamiltonian* as

$$\check{H}(t, z, p, \mathbf{A}) = \sup_{\bar{u} \in \mathcal{U}} \left[ -B(t, z, \bar{u}) \cdot p - \frac{1}{2} \text{tr } \mathcal{A}(t, z, \bar{u}) \mathbf{A} - g(t, z, \bar{u}) \right]$$



For the finite-cost finite-horizon problem, the HJB PDE is then defined as

$$-\frac{\partial v}{\partial t} + \check{H}(t, z, \nabla v, D_x^2 v) = 0, \quad (t, z) \in \mathcal{Q},$$

where  $D_x^2$  is the Hessian operator. Let

$$\partial^* \mathcal{Q} = ([t_0, t_1] \times \partial \mathcal{O}) \cup (\{t_1\} \times \mathcal{O}),$$

denote the boundary of the product state space. Then, the boundary conditions for the HJB PDE are given by

$$v(t, z) = \psi(t, z) \quad \text{for } (t, z) \in \partial^* \mathcal{Q}.$$

This PDE is of parabolic type. It is called uniformly parabolic if there exists a  $c > 0$  such that

$$\sum_{i,j}^d \mathcal{A}[i, j](t, x, \bar{u}) \xi_i \xi_j \geq c |\xi|^2, \quad (3)$$

for  $\xi \in \mathbb{R}^d$ , otherwise it is of degenerate parabolic type [19].

For discounted-cost infinite-horizon problems, the Hamiltonian and corresponding HJB PDE are no longer time dependent. The HJB PDE is then defined as

$$\beta v + \check{H}(z, \nabla v, D_x^2 v) = 0, \quad z \in \mathcal{O},$$

with boundary conditions

$$v(z) = \psi(z), \quad z \in \partial \mathcal{O}.$$

If the diffusion term satisfies the condition in Equation (3) and  $\mathcal{O}$  is bounded, then the HJB PDE is a uniformly elliptic partial differential equation having unique solution that is smooth. Otherwise, the equation is of degenerate elliptic type [19].

### III. THE MARKOV CHAIN APPROXIMATION METHOD

In this section, we provide background for a solution method based on discretization that forms the basis of our computational framework. The Markov chain approximation (MCA) [39] and similar methods, e.g., the method prescribed by Tsitsiklis [60], for solving the stochastic optimal control problem rely on first discretizing the state space and dynamics described by Equation (1) and then solving the resulting discrete-time and discrete-space Markov Decision Process (MDP). The discrete MDP can be solved using standard techniques such as Value Iteration (VI) or Policy Iteration (PI) or other approximate dynamic programming techniques [3], [6], [7], [39], [51].

In Section III-A, we provide a brief overview of discrete MDPs. In Section III-B we describe the Markov chain approximation method for discretizing continuous stochastic optimal control problems. Finally, in Section III-C, we describe three standard algorithms to solve discrete MDPs, namely value iteration, policy iteration, as well as multilevel methods.

#### A. Discrete-time discrete-space Markov decision processes

The Markov chain approximation method relies on discretizing the state space of the underlying stochastic dynamical system, for instance, using a grid. The discretization is parametrized by the discretization step, denoted by  $h \in \mathbb{R}_+$ , of the grid. The discretization is finer for smaller values of  $h$ .

The MDP resulting from the Markov chain approximation method with a discretization step  $h$  is a tuple, denoted by  $\mathcal{M}^h = (\mathcal{X}^h, \mathcal{U}, p^h, g^h, \psi^h)$ , where  $\mathcal{X}^h$  is the set of discrete states,  $p^h(\cdot, \cdot, \cdot) : \mathcal{X}^h \times \mathcal{X}^h \times \mathcal{U} \rightarrow [0, 1]$  is function that denotes the transition probabilities satisfying  $\sum_{z' \in \mathcal{X}^h} p^h(z, z' | \bar{u}) = 1$  for all  $z \in \mathcal{X}^h$  and all  $\bar{u} \in \mathcal{U}$ ,  $g^h$  is the stage cost of the discrete system, and  $\psi^h$  is the terminal cost of the discrete system.

The transition probabilities replace the drift and diffusion terms of the stochastic dynamical system as the description for the evolution of the state. For example, when the process is at state  $z \in \mathcal{X}^h$  and action  $\bar{u} \in \mathcal{U}$  is applied, the next state of the process becomes  $z' \in \mathcal{X}^h$  with probability  $p^h(z, z' | \bar{u})$ .

In this discrete setting, Markov policies are now mappings  $\mu^h : \mathcal{T} \times \mathcal{X}^h \rightarrow \mathcal{U}$  defined from a discrete state space rather than from the continuous space  $\mathcal{X}$ . Furthermore, the cost functional becomes a multidimensional array  $c_{\mu^h} : \mathcal{T} \times \mathcal{X}^h \rightarrow \mathbb{R}$ . The cost associated with a particular trajectory and policy, for a finite horizon, discrete time system, i.e.,  $t_0 = 0, t_1 = 1, t_2 = 2, \dots$ , can be written as

$$c_{\mu^h}(t_0, z) = \mathbb{E} \left[ \sum_{i=1}^N \gamma^i g^h(t_k, x(t_k), \mu^h(t_k, x(t_k))) + \psi^h(t_N, x(t_N)) \right], \quad x(t_0) = z$$

for  $z \in \mathcal{X}^h$ , where  $0 < \gamma < 1$  is the discount factor. The *Bellman equation*, a discrete analogue of the HJB equation, describing the optimality of this discretized problem can then be written as

$$v^h(t_k, z) = \min_{\bar{u} \in \mathcal{U}} \left[ g^h(t_k, z, \bar{u}) + \gamma \sum_{z' \in \mathcal{X}^h} p^h(z, z' | \bar{u}) v^h(t_{k+1}, z') \right],$$

where  $v^h$  is the optimal discretized value function and satisfies the Bellman equation

$$v^h(t, z) = \inf_{\mu^h} c_{\mu^h}(t, z).$$

Therefore, it also solves the discrete MDP [6].

In this paper, we will consider the value function to be only a function of the state and independent of time. In this case, the stage costs are also assumed to be independent of time, i.e.,  $g^h : \mathcal{X}^h \rightarrow \mathbb{R}$ , and the Bellman equation is written as

$$v^h(z) = \min_{\bar{u} \in \mathcal{U}} \left[ g^h(z, \bar{u}) + \gamma \sum_{z' \in \mathcal{X}^h} p^h(z, z' | \bar{u}) v^h(z') \right]. \quad (4)$$

Analytically, this assumption is without loss of any generality. If the value function is time dependent, for example in finite horizon problems, one can always augment the state space with the time variable to create a new value function that is a function of the augmented (or product) state space.

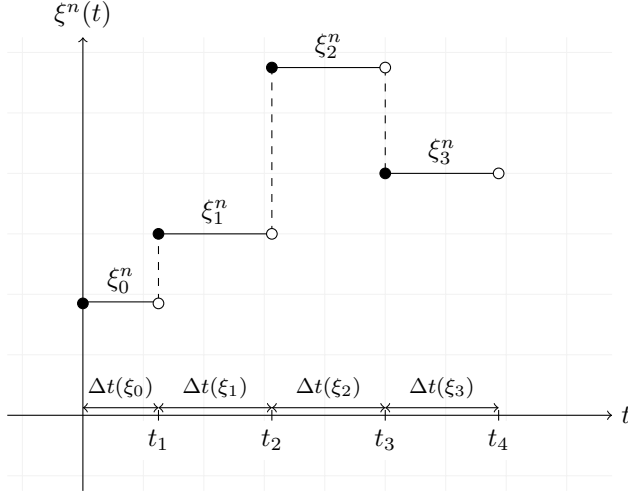


Fig. 1: An illustration of a continuous-time interpolation of a discrete process arising from the Markov chain approximation.

### B. Markov chain approximation method

The MCA method, developed by Kushner and co-workers [37]–[39], constructs a sequence of discrete MDPs such that the solution of the MDPs converge to the solution of the original continuous-time continuous-space problem.

Let  $\{\mathcal{M}^{h_\ell} : \ell \in \mathbb{N}\}$  be a sequence of MDPs, where each  $\mathcal{M}^{h_\ell} = (\mathcal{X}^{h_\ell}, \mathcal{U}, p^{h_\ell}, g^{h_\ell}, \psi^{h_\ell})$  is defined as before. Define  $\partial\mathcal{X}^{h_\ell}$  as the subset of  $\mathcal{X}^{h_\ell}$  that falls on the boundary of  $\mathcal{X}$ , i.e.,  $\partial\mathcal{X}^{h_\ell} = \partial\mathcal{X} \cap \mathcal{X}^{h_\ell}$ . Let  $\{\Delta t^\ell : \ell \in \mathbb{N}\}$ , where  $\Delta t^\ell : \mathcal{X}^{h_\ell} \rightarrow \mathbb{R}_+$ , be a sequence holding times. Let  $\{\xi_i^n : i \in \mathbb{N}\}$ , where  $\xi_i^n \in \mathcal{X}^{h_\ell}$ , be a (random) sequence of states that describe the trajectory of  $\mathcal{M}^{h_\ell}$ . We use holding times as interpolation intervals to generate a continuous-time trajectory from this discrete trajectory as follows. With a slight abuse of notation, let  $\xi^n : \mathbb{R}_{\geq 0} \rightarrow \mathcal{X}^{h_\ell}$  denote the continuous-time function defined as follows:  $\xi^n(\tau) = \xi_i^n$  for all  $\tau \in [t_i^\ell, t_{i+1}^\ell)$ , where  $t_i^\ell = \sum_{k=0}^{i-1} \Delta t^\ell(\xi_k)$ . Let  $\{u_i^\ell : i \in \mathbb{N}\}$ , where  $u_i^\ell \in \mathcal{U}$ , be a sequence of control inputs defined for all  $\ell \in \mathbb{N}$ . Then, we define the continuous time interpolation of  $\{u_i^\ell : i \in \mathbb{N}\}$  as  $u^\ell(\tau) = u_i^\ell$  for all  $\tau \in [t_i^\ell, t_{i+1}^\ell)$ . An illustration of this interpolation is provided in Figure 1.

The following result by Kushner and co-workers characterizes the conditions under which the trajectories and value functions of the discrete MDPs converge to those of the original continuous-time continuous-space stochastic system.

**Theorem 1** (See Theorem 10.4.1 by Kushner and Dupuis [39]). *Suppose the sequence  $\{\mathcal{M}^{h_\ell} : \ell \in \mathbb{N}\}$  of MDPs and the sequence  $\{\Delta t^\ell : \ell \in \mathbb{N}\}$  holding times satisfy the following conditions: For any sequence of inputs  $\{u_i^\ell : i \in \mathbb{N}\}$  and the resulting sequence of trajectories  $\{\xi_i^\ell : i \in \mathbb{N}\}$  if*

$$\lim_{\ell \rightarrow \infty} \Delta t^\ell(z) = 0, \text{ for all } z \in \mathcal{X},$$

and

$$\lim_{\ell \rightarrow \infty} \frac{\mathbb{E}[\xi_{i+1}^\ell - \xi_i^\ell | \xi_i^\ell = z, u_i^\ell = \bar{u}]}{\Delta t^\ell(z)} = B(z, \bar{u}),$$

$$\lim_{\ell \rightarrow \infty} \frac{\text{Cov}[\xi_{i+1}^\ell - \xi_i^\ell | \xi_i^\ell = z, u_i^\ell = \bar{u}]}{\Delta t^\ell(z)} = \mathcal{D}(z, \bar{u}),$$

for all  $z \in \mathcal{X}$  and  $\bar{u} \in \mathcal{U}$ . Then, the sequence  $\{(\xi^\ell, u^\ell) : \ell \in \mathbb{N}\}$  of interpolations converges in distribution to  $(x, u)$  that solves the integral equation with differential form given by (1). Let  $v^{h_\ell}$  denote the optimal value function for the MDP  $\mathcal{M}^{h_\ell}$ . Then, for all  $z \in \mathcal{X}^{h_\ell}$ ,

$$\lim_{\ell \rightarrow \infty} |v^{h_\ell}(z) - v(z)| = 0.$$

The conditions of this theorem are called *local consistency conditions*. Roughly speaking, the theorem states that the trajectories of the discrete MDPs will converge to the trajectories of the original continuous-time stochastic dynamical system if the local consistency conditions are satisfied. Furthermore, in that case, the value function of the discrete MDPs also converge to that of the original stochastic optimal control problem. A discretization that satisfies the local consistency conditions is called a *consistent discretization*. Once a consistent discretization is obtained, standard dynamic programming algorithms such as value iteration or policy iteration [3] can be used for its solution.

1) *Discretization procedures*: In this section, we provide a general discretization framework described by Kushner and Dupuis [39] along with a specific example. For the rest of the paper, we will drop the subscript  $\ell$  from  $h_\ell$  for the sake of brevity, and simply refer to the discretization step with  $h$ .

Let  $\mathcal{X}^h \subset \mathcal{X}$  denote a discrete set of states. For each state  $z \in \mathcal{X}^h$  define a finite set of vectors  $M(z) = \{v_{i,z} : i < m(z)\}$ , where  $v_{i,z} \in \mathbb{R}^d$  and  $m(z) : \mathcal{X}^h \rightarrow \mathbb{N}$  is uniformly bounded. These vectors denote directions from a state  $z$  to a neighboring set of states  $\{y : y = z + hv_{i,z}, i \leq m(z)\} \subset \mathcal{X}^h$ . A valid discretization is described by the functions  $q_i^1(z) : \mathcal{X}^h \rightarrow \mathbb{R}$  and  $q_i^0(z, \bar{u}) : \mathcal{X}^h \times \mathcal{U} \rightarrow \mathbb{R}$  that satisfy

$$B(z, \bar{u}) = \sum_{v_{i,z} \in M(z)} q_i^0(z, \bar{u}) v_{i,z}, \text{ for all } \bar{u},$$

$$\mathcal{D}(z) = \sum_{v_{i,z} \in M(z)} q_i^1(z) v_{i,z} v_{i,z}',$$

$$\sum_{v_{i,z} \in M(z)} q_i^1(z) v_{i,z} = 0,$$

$$hq_i^0(z, \bar{u}) + q_i^1(z) \geq 0,$$

$$q_i^1(z) > 0,$$

where the third condition guarantees that  $q_i^1$  only contribute to the variance of the chain and not the mean, and the fourth and fifth conditions guarantee non-negative transition probabilities, which we verify momentarily.

After finding  $q_i^1$  and  $q_i^0$  that satisfy these conditions, we are ready to define the approximating MDP  $\mathcal{M}^h = (\mathcal{X}^h, \mathcal{U}, p^h, g^h, \psi^h)$ . First, define the normalizing constant

$$q^h(z, \bar{u}) = \sum_{v_{i,z} \in M(z)} [hq_i^0(z, \bar{u}) + q_i^1(z)]$$

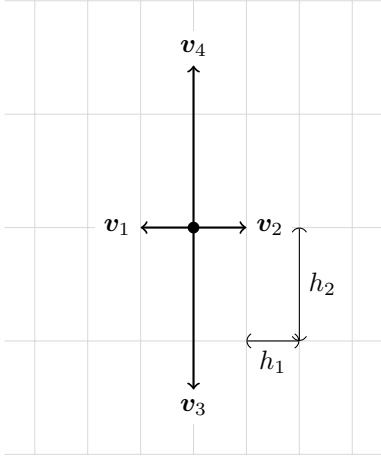


Fig. 2: Sample discretization of a two-dimensional state space.

Then, the discrete MDP is defined by the state space  $\mathcal{X}^{h_\ell}$ , the control space  $\mathcal{U}$ , the transition probabilities

$$p^h(z, z + h\mathbf{v}_{i,z}|\bar{u}) = \frac{hq_i^0(z, \bar{u}) + q_i^1(z)}{q^h(z, \bar{u})},$$

and the stage costs

$$g^h(z, \bar{u}) = \frac{\Delta t^h(z, \bar{u})}{q^h(z, \bar{u})} g(z, \bar{u}).$$

The discount factor is

$$\gamma = \exp(-\beta \Delta t^h).$$

Finally, define the interpolation interval

$$\Delta t^h(z, \bar{u}) = \frac{h^2}{q^h(z, \bar{u})}.$$

These conditions satisfy local consistency [39].

2) *Upwind differencing*: One realization of the framework described above is generated based on upwind differencing. This procedure tries to “push” the current state of the system in the direction of the drift dynamics on average, and we describe this method here and use it for all of the numerical examples in Section VII.

Consider the sample discretization of a two dimensional state space provided in Figure 2. The state is discretized with a spacing of  $h_1$  in the first dimension and  $h_2$  in the second dimension. In this example, the directions  $\mathbf{v}_{i,z}$  are independent of  $z$ , and thus we drop the subscript  $z$  and refer simply to  $\mathbf{v}_i$ . Let the discretization step be defined according to  $h = \min(h_1, h_2)$  so that transition directions  $\mathbf{v}_i$  are defined as

$$\mathbf{v}_1 = -\frac{h_1}{h}\mathbf{e}_1, \quad \mathbf{v}_2 = \frac{h_1}{h}\mathbf{e}_1, \quad \mathbf{v}_3 = -\frac{h_2}{h}\mathbf{e}_2, \quad \mathbf{v}_4 = \frac{h_2}{h}\mathbf{e}_2.$$

This is a two dimensional state space so that the output of the drift can be indexed according to  $B[1]$  and  $B[2]$ . Similarly, the diffusion is defined such that  $\mathcal{A}(z) = \mathcal{D}(z)\mathcal{D}^T(z)$  be a diagonal matrix such that  $\mathcal{A}(z) = \text{diag}([\mathcal{A}[1, 1](z), \mathcal{A}[2, 2](z)])$ , let  $B[i](z, \bar{u})^+ = \max(0, B[i](z, \bar{u}))$ , and finally let  $B[i](z, \bar{u})^- = \max(0, -B[i](z, \bar{u}))$ . The directions  $\mathbf{v}_i$  do not need to be the same length as the discretization size of the

state space. In particular, their length is the ratio between the grid size in the particular direction and the finest discretization. This characteristic will ensure local consistency even for cases when the discretization is different in each direction by causing the transition probabilities to take the relative discretizations of each dimension into account. Then, a locally consistent discretization is one defined by

$$\begin{aligned} q_1^0(z, \bar{u}) &= \frac{h}{h_1} B[1](z, \bar{u})^-, & q_0^1 &= \left(\frac{h}{h_1}\right)^2 \frac{\mathcal{A}[1, 1](z)^2}{2}, \\ q_2^0(z, \bar{u}) &= \frac{h}{h_1} B[1](z, \bar{u})^+, & q_1^1 &= \left(\frac{h}{h_1}\right)^2 \frac{\mathcal{A}[1, 1](z)^2}{2}, \\ q_3^0(z, \bar{u}) &= \frac{h}{h_2} B[2](z, \bar{u})^-, & q_2^1 &= \left(\frac{h}{h_2}\right)^2 \frac{\mathcal{A}[2, 2](z)^2}{2}, \\ q_4^0(z, \bar{u}) &= \frac{h}{h_2} B[2](z, \bar{u})^+, & q_3^1 &= \left(\frac{h}{h_2}\right)^2 \frac{\mathcal{A}[2, 2](z)^2}{2}, \end{aligned}$$

Suppose the drift is defined such that  $B[1](z, \bar{u}) > 0$  while  $B[2](z, \bar{u}) < 0$  then

$$\begin{aligned} \sum_{i \in M(z)} q_i^0(z, \bar{u}) \mathbf{v}_i &= \underbrace{\frac{h}{h_1} B[1](z, \bar{u})^+}_{q_2^0(z, \bar{u})} \underbrace{\frac{h_1}{h} \mathbf{e}_1}_{\mathbf{v}_2} \\ &\quad + \underbrace{\frac{h}{h_2} B[2](z, \bar{u})^-}_{q_3^0(z, \bar{u})} \underbrace{\left(-\frac{h_2}{h} \mathbf{e}_2\right)}_{\mathbf{v}_3} \\ &= \begin{bmatrix} B[1](z, \bar{u}) \\ B[2](z, \bar{u}) \end{bmatrix} = B(z, \bar{u}) \end{aligned}$$

All the other conditions can be similarly verified.

We can also analyze the computational cost of this upwind differencing procedure. The computation of the transition probabilities, for some state  $z$  and control  $\bar{u}$ , requires the evaluation of the drift and diffusion. Suppose that this evaluation requires  $n_{\text{op}}$  operations. Assembling each  $q_i^0(z, \bar{u})$  and  $q_i^1$  requires two operations: multiplication and division. Since there are  $2d$  neighbors for each  $z$ , the evaluation of all of them requires  $4d$  operations. Next, the computation of the normalization  $q^h$  involves summing all of the  $q_i^0$  and  $q_i^1$ , a procedure requiring  $4d$  operations. Computing the interpolation interval requires a single division, computing the discrete stage cost requires a division and multiplication, and computing the discount factor requires exponentiation. Together, these operations mean that the computational complexity of discretizing the SOC for some state  $z$  and control  $\bar{u}$  using upwind differencing is linear with dimension

$$\mathcal{O}(n_{\text{op}} + d). \quad (5)$$

3) *Boundary conditions*: The discretization methods described in the previous section apply to the interior nodes of the state space. In order to numerically solve optimal stochastic control problems, however, one typically needs to limit the state space to a particular region. In order to utilize low-rank tensor based methods in high dimensions, we design  $\mathcal{O}$  to be a hypercube. Due to this state truncation we are required to assign boundary conditions for the discrete Markov process. Three boundary conditions are commonly used: periodic, absorbing, and reflecting boundary conditions.

A *periodic boundary condition* maps one side of the domain to the other. For example consider  $\mathcal{O} = (-1, 1)^2$ . Then, if we define a periodic boundary condition for the first dimension, we mean that  $z = (-1, \cdot)$  and  $z' = (1, \cdot)$  are equivalent states.

An *absorbing boundary condition* dictates that if the Markov process enters  $\partial\mathcal{O}$  at the exit time  $\tau$ , then the process terminates and terminal costs are incurred.

A *reflecting boundary condition* is often imposed when one does not want to end the process at the boundary and periodic boundaries are not appropriate. In this case, the stochastic process is modeled with a jump diffusion. The jump diffusion term is responsible for keeping the process within  $\mathcal{O}$ . In our case, we will assume that the jump diffusion term instantaneously “reflects” the process using an orthogonal projection back into the state space  $\mathcal{O}$ . For example, if the system state is  $z \in \mathcal{O}$  and the Markov process transitions to  $z' = z + he_k$  such that  $z' \in \partial\mathcal{O}$ , then the system immediately returns to the state  $z$ . Therefore, we can eliminate  $z'$  from the discretized state space and adjust the self transition probability to be

$$p^h(z, z|\bar{u}) \leftarrow p^h(z, z|\bar{u}) + p^h(z, z'|\bar{u}).$$

In other words, the probability of self transitioning is increased by the probability of transitioning to the boundary.

### C. Value iteration, policy iteration, and multilevel methods

In this section, we describe algorithms for solving the discounted-cost infinite-horizon MDP given by Equation (4). In particular, we describe the value iteration (VI) algorithm and the policy iteration (PI) algorithm. Then, we describe a multi-level algorithm that is able to use coarse-grid solutions to generate solutions of fine-grid problems. FT-based versions of these algorithms will then be described in Section V.

1) *DP equations*: Let  $\mathcal{R}^h$  be the set of real-valued functions  $w^h : \mathcal{X}^h \rightarrow \mathbb{R}$ . Define the functional  $H^h : \mathcal{X}^h \times \mathcal{U} \times \mathcal{R}^h$  as

$$H^h(z, \bar{u}, w^h) := g^h(z, \bar{u}) + \gamma \sum_{z' \in \mathcal{X}^h} p^h(z, z'|\bar{u}) w^h(z'). \quad (6)$$

For a given policy  $\mu$ , define operator  $T_\mu^h : \mathcal{R}^h \rightarrow \mathcal{R}^h$  as

$$T_\mu^h(w^h)(z) := H^h(z, \mu(z), w^h), \quad \forall z \in \mathcal{X}^h, w^h \in \mathcal{R}^h$$

Define the mapping  $T^h : \mathcal{R}^h \rightarrow \mathcal{R}^h$ , which corresponds to the Bellman equation given by Equation (4), as

$$T^h(w^h)(z) := \min_{\mu(z)} H^h(z, \bar{u}, w^h), \quad \forall z \in \mathcal{X}^h, w^h \in \mathcal{R}^h.$$

Using these operators we can denote two important fixed-point equations. The first describes the value function  $w^h$  that corresponds to a fixed policy  $\mu$

$$w^h = T_\mu^h(w^h). \quad (7)$$

The second equation describes the optimal value function  $v^h$

$$v^h = T^h(v^h) \quad (8)$$

These equations are known as the dynamic programming equations.

2) *Assumptions for convergence*: Three assumptions are required to guarantee existence and uniqueness of the solution to the dynamic programming equations and to validate the convergence of their associated solution algorithms.

**Assumption 1** (Assumption A1.1 by Kushner and Dupuis [39]). *The functions  $p^h(z, z'|\bar{u})$  and  $g^h(z, \bar{u})$  are continuous functions of  $\bar{u}$  for all  $z, z' \in \mathcal{X}^h$ .*

The second assumption involves contraction.

**Definition 1** (Contraction). *Let  $\mathcal{Y}$  be a normed vector space with the norm  $\|\cdot\|$ . A function  $f : \mathcal{Y} \rightarrow \mathcal{Y}$  is a contraction mapping if for some  $\gamma \in (0, 1)$  we have*

$$\|f(y) - f(y')\| \leq \gamma \|y - y'\|, \quad \forall y, y' \in \mathcal{Y}.$$

**Assumption 2** (Assumption A1.2 by Kushner and Dupuis [39]). *(i) There is at least one admissible feedback policy  $\mu$  such that  $T_\mu^h$  is a contraction, and the infima of the costs over all admissible policies is bounded from below. (ii)  $T_\mu^h$  is a contraction for any feedback policy for which the associated cost is bounded.*

The third assumption involves the repeated application  $T_\mu^h$ .

**Assumption 3** (Assumption A1.3 by Kushner and Dupuis [39]). *Let  $\mathbf{P}_\mu = \{p^h(z, z'|\mu(z)) : z, z' \in \mathcal{X}^h\}$  be the matrix formed by the transition probabilities of the discrete-state MDP for a fixed policy  $\mu$ . If the value functions associated with the use of policies  $\mu_1, \dots, \mu_n, \dots$  in sequence, is bounded, then*

$$\lim_{n \rightarrow \infty} \mathbf{P}_{\mu_1} \mathbf{P}_{\mu_2} \cdots \mathbf{P}_{\mu_n} = 0$$

3) *Value iteration algorithm*: The VI algorithm is a fixed-point (FP) iteration aimed at computing the optimal value function  $v^h$ . It works by starting with an initial guess  $v_0^h \in \mathcal{R}^h$  and defining a sequence of value functions  $\{v_k^h\}$  through the iteration  $v_{k+1}^h = T^h(v_k^h)$ . Theorem 2 guarantees the convergence of this algorithm under certain conditions.

**Theorem 2** (Jacobi iteration, Theorem 6.2.2 by Kushner and Dupuis [39]). *Let  $\mu$  be an admissible policy such that  $T_\mu^h$  is a contraction. Then for any initial vector  $w_0^h \in \mathcal{R}^h$ , the sequence  $w_k^h$  defined by*

$$w_{k+1}^h = T_\mu^h(w_k^h) \quad (9)$$

*converges to  $w^h$ , the unique solution to Equation (7). Assume Assumptions 1, 2, and 3. Then for any vector  $v_0^h \in \mathcal{R}^h$ , the sequence recursively defined by*

$$v_{k+1}^h = T^h(v_k^h) \quad (10)$$

*converges to the optimal value function  $v^h$ , the unique solution to Equation (8).*

Indeed, Equation (10) is the fixed-point iteration that is the value iteration algorithm. In Section V-B, we will describe an FT-based version of this algorithm.

4) *Policy iteration algorithm*: Policy iteration (PI) is another method for solving discrete MDPs. Roughly, it is analogous to a gradient descent method, and our experiments indicate that it generally converges faster than the VI algorithm. The basic idea is to start with a Markov policy  $\mu_0$  and to generate a sequence of policies  $\{\mu_k\}$  according to

$$\mu_k = \arg \min_{\mu} [T_{\mu}^h(w_{k-1}^h)] \quad (11)$$

The resulting value functions  $\{w_k^h\}$  are solutions of Equation (7), i.e.,

$$w_k^h = T_{\mu_k}^h(w_k^h). \quad (12)$$

Theorem 3 provides the conditions under which this iteration converges.

**Theorem 3** (Policy iteration, Theorem 6.2.1 by Kushner and Dupuis [39]). *Assume Assumptions 1 and 2. Then there is a unique solution to Equation (8), and it is the infimum of the value functions over all time independent feedback policies. Let  $\mu_0$  be an admissible feedback policy such that the corresponding value function  $w_0^h$  is bounded. For  $k \geq 1$ , define the sequence of feedback policies  $\mu_k$  and costs  $w_k^h$  recursively by Equation (11) and Equation (12). Then  $w_k^h \rightarrow v^h$ . Under the additional condition given by Assumption 3,  $v^h$  is the infimum of the value functions over all admissible policies.*

Note that policy iteration requires the solution of a linear system in Equation (12). Furthermore, Theorem 2 states that since  $T_{\mu}^h$  is a contraction mapping that a fixed-point iteration can also be used to solve this system. This property leads to a modification of the policy iteration algorithm called *optimistic policy iteration* [6], which substitutes  $n_{fp}$  steps of fixed-point iterations for solving (12) to create a computationally more efficient algorithm. The assumptions necessary for convergence of optimistic PI are the same as those for PI and VI. Details are given by Bertsekas [6]. In Section V-C, we will describe an FT-based version of optimistic PI algorithm.

5) *Multilevel algorithms*: Multigrid techniques [9], [59] have been successful at obtaining solutions to many problems by exploiting multiscale structure of the problem. For example, they are able to leverage solutions of linear systems at coarse discretization levels for solving finely discretized systems.

We describe how to apply these ideas within dynamic programming framework for two purposes: the initialization of fine-grid solutions with coarse-grid solutions and for the solution of the linear system in Equation (12) within policy iteration. Our experiments indicate that fine-grid problems typically require more iterations to converge, and initialization with a coarse-grid solution offers dramatic computational gains. Since we expect the solution to converge to the continuous solution as the grid is refined, we expect the number of iterations required for convergence to decrease as the grid is refined.

The simplest multi-level algorithm is the one-way discretization algorithm that sequentially refines coarse-grid solutions of Equation (8) by searching for solutions on a grid starting from an initial guess obtained from the solution of a coarser problem. This procedure was analyzed in detail for shortest path or MDP style problems by Chow and Tsitsiklis

in [13]. The pseudocode for this algorithm provided in Algorithm 1. In Algorithm 1, a set of  $\kappa$  discretization levels  $\{h_1, h_2, \dots, h_{\kappa}\}$  such that for  $j > i$  we have  $h_j > h_i$ , are specified. Furthermore, the operator  $I_{2h}^h$  interpolates the solution of the  $h_{k+1}$  grid onto the  $h_k$  grid. Then a sequence of problems, starting with the coarsest, are solved until the fine-grid solution is obtained.

---

**Algorithm 1** One-way multigrid [13]

---

**Require:** Set of discretization levels  $\{h_1, h_2, \dots, h_{\kappa}\}$ ; Initial cost function  $v_0^{h_{\kappa}}$

- 1:  $v^{h_{\kappa}} \leftarrow$  Solve Equation (8) starting from  $v_0^{h_{\kappa}}$
- 2: **for**  $k = \kappa - 1 \dots 1$  **do**
- 3:  $v_0^{h_k} = I_{2h}^h v^{h_{k+1}}$
- 4:  $v^{h_k} \leftarrow$  Solve Equation (8) starting from  $v_0^{h_k}$
- 5: **end for**

---

Multigrid techniques can also be used for solving the linear system in Equation (7) within the context of policy iteration. Recall that for a fixed policy  $\mu$ , this system can be equivalently written using a linear operator  $\Pi_{\mu}^h$  defined according to

$$\Pi_{\mu}^h(w^h)(z) \equiv w^h(z) - \gamma \sum_{z' \in \mathcal{X}^h} p^h(z, z' | \mu(z)) w^h(z'), \quad \forall z \in \mathcal{X}^h.$$

Therefore, for a fixed policy  $\mu$  the corresponding value function satisfies

$$\Pi_{\mu}^h(w^h) = g^h. \quad (13)$$

Typically, we do not expect to satisfy this equation exactly, rather we will have an approximation  $\hat{w}^h$  that yields a non-zero residual

$$r^h = g^h - \Pi_{\mu}^h(\hat{w}^h). \quad (14)$$

In addition to the residual, we can define the difference between the approximation and the true minimum as  $\Delta w^h = w^h - \hat{w}^h$ . Since,  $\Pi_{\mu}^h$  is a linear operator, we can replace  $\hat{w}^h$  in Equation (14) to obtain

$$\begin{aligned} r^h &= g^h - \Pi_{\mu}^h(w^h - \Delta w^h) \\ &= g^h - \Pi_{\mu}^h(w^h) + \Pi_{\mu}^h(\Delta w^h) \\ &= \Pi_{\mu}^h(\Delta w^h) \end{aligned}$$

Thus, if solve for  $\Delta w^h$ , then we can update  $\hat{w}^h$  to obtain the solution

$$w^h = \Delta w^h + \hat{w}^h.$$

In order for multigrid to be a successful strategy, we typically assume that the residual  $r^h$  is “smooth,” and therefore we can potentially solve for  $\Delta w^h$  on a coarser grid. The coarse grid residual is

$$r^{2h} = \Pi_{\mu}^{2h}(\Delta w^{2h})$$

where we now choose the residual to be the restriction, denoted by operator  $I_h^{2h}$ , of the fine-grid residual

$$r^{2h} = I_h^{2h} r^h.$$

Combining these two equations we obtain an equation for  $\Delta w^{2h}$

$$\Pi_{\mu}^{2h}(\Delta w^{2h}) = I_h^{2h} r^h \quad (15)$$

Note that the relationship between the linear operators  $T_\mu^{2h}$  and  $\Pi_\mu^{2h}$  displayed by Equations (7) and (13) lead to an equivalent equation for the error given by

$$\Delta w^{2h} = T_\mu^{2h}(\Delta w^{2h}; r^{2h}), \quad (16)$$

where we specifically denote that the stage cost is replaced by  $r^{2h}$ . Since  $T_\mu^{2h}$  is a contraction mapping we can use the fixed-point iteration in Equation (9) to solve this equation.

After solving the system in Equation (15) above we can obtain the correction at the fine grid

$$\hat{w}^h \leftarrow \hat{w}^h + I_{2h}^h \Delta w^{2h} \quad (17)$$

In order, to obtain smooth out the high frequency components of the residual  $r^h$  one must perform “smoothing” iteration instead of the typical iteration  $T_\mu^h$ . These iterations are typically Gauss-Seidel relaxations or weighted Jacobi iterations. Suppose that we start with  $\hat{w}_k^h$ , then using the weighted Jacobi iteration we obtain an update  $\hat{w}_{k+1}^h$  through the following two equations

$$\begin{aligned} \tilde{w}^h &= T_\mu^h(\hat{w}_k^h, g^h), \\ \hat{w}_{k+1}^h &= \omega \tilde{w}^h + (1 - \omega) \hat{w}_k^h, \end{aligned}$$

for  $\omega > 1$ . To shorten notation, we will denote these equations by the operator  $T_{c,\mu}^h$  such that

$$\hat{w}_{k+1}^h = T_{c,\mu}^h(\hat{w}_k^h, g^h).$$

We have chosen to use the weighted Jacobi iteration since it can be performed by treating the linear operator  $T_\mu^h$  as a black-box FP iteration, i.e., the algorithm takes as input a value function and outputs another value function. Thus, we can still wrap the low-rank approximation scheme around this operator. A relaxed Gauss-Seidel relaxation would require sequentially updating elements of  $\hat{w}_k^h$ , and then using these updated elements for other element updates. Details on the reasons for these smoothing iterations within multigrid is available in the existing literature [9], [39], [59].

Combining all of these notions we can design many multigrid methods. We will introduce an FT-based version of the two-level V-grid in Algorithm 5 in the next section.

#### IV. LOW-RANK COMPRESSION OF FUNCTIONS

The Markov chain approximation method is often computationally intractable for problems instances with state spaces embedded in more than a few dimensions. This curse of dimensionality, or exponential growth in storage and computation complexity, arises due to state-space discretization. To mitigate the curse of dimensionality, we believe that algorithms for solving general stochastic optimal control problems must be able to

- 1) exploit function structure to perform compression with polynomial time complexity, and
- 2) perform multilinear algebra with functions in compressed format in polynomial time.

The first capability ensures that value functions can be *represented* on computing hardware. The second ensures that they can be used for computation.

In this section, we describe a method for “low-rank” representation of multivariate functions, and algorithms that allow us to perform multilinear algebra in this representation. The algorithms presented in this section were introduced in earlier work by the authors [24]. In that paper, algorithms for approximating multivariate black box functions and performing computations with the resulting approximation are described. We review low-rank function representations in Section IV-A and the approximation algorithms in Section IV-B.

##### A. Low-rank function representations

Let  $\mathcal{X}$  be a tensor product of closed intervals  $\mathcal{X} := [a_1, b_1] \times [a_2, b_2] \times \cdots \times [a_d, b_d]$ , with  $a_i, b_i \in \mathbb{R}$  and  $a_i < b_i$  for  $i = 1, \dots, d$ . A *low-rank* approximation of a function  $f : \mathcal{X} \rightarrow \mathbb{R}$  represents the function’s values as the sum of a few products of univariate functions  $f_j^{(i)} : [a_i, b_i] \rightarrow \mathbb{R}$ , i.e.,

$$f(x_1, \dots, x_d) \approx \sum_{j=1}^R f_j^{(1)}(x_1) \cdots f_j^{(d)}(x_d).$$

This approximation is called the canonical polyadic (CP) decomposition [12]. The CP format is defined by sets of univariate functions  $\mathcal{F}_i^{CP} = \{f_1^{(i)}(x_i), \dots, f_R^{(i)}(x_i)\}$  for  $i = 1, \dots, d$ . The storage complexity of the CP format is clearly linear with dimension, but also depends on the storage complexity of each  $f_j^{(i)}$ . For instance, if each input dimension is discretized into  $n$  grid points, so that each univariate function is represented by  $n$  values, then the storage complexity of the CP format is  $\mathcal{O}(dnR)$ . This complexity is linear with dimension, linear with discretization level, and linear with rank  $R$ . Thus, for a class of functions whose ranks grow polynomially with dimension, i.e.,  $R(d) = \mathcal{O}(d^p)$  for some  $p \in \mathbb{N}$ , polynomial storage complexity is attained.

Contrast this with the representation of  $f(x_1, \dots, x_d)$  as a lookup table. If the lookup table is obtained by discretizing each input variable into  $n$  points, the storage requirement is  $\mathcal{O}(n^d)$ , which grows exponentially with dimension.

Regardless of the representation of each  $f_j^{(i)}$ , the complexity of this representation is always *linear* with dimension. Intuitively, for approximately separable functions, storing many univariate functions requires fewer resources than storing a multivariate function. In the context of stochastic optimal control, the CP format has been used for the special case when the control is unconstrained and the dynamics are affine with control input [29].

Using the canonical decomposition can be problematic in practice because the problem of determining the canonical rank of a discretized function, or tensor, is NP complete [28], [36], and the problem of finding the best approximation in Frobenius norm for a given rank can be ill-posed [15]. Instead of the canonical decomposition, we propose using a continuous variant of the *tensor-train* (TT) decomposition [24], [47], [49] called the *function-train* (FT) [24]. In these formats, the best fixed-rank approximation problem is well posed, and the approximation can be computed using a sequence of matrix factorizations.

A multivariate function  $f : \mathcal{X} \rightarrow \mathbb{R}$  in FT format is represented as

$$f(x_1, \dots, x_d) = \sum_{\alpha_0=1}^{r_0} \dots \sum_{\alpha_d=1}^{r_d} f_{\alpha_0, \alpha_1}^{(1)}(x_1) \dots f_{\alpha_{d-1}, \alpha_d}^{(d)}(x_d), \quad (18)$$

where  $r_i \in \mathbb{Z}_+$  are the FT ranks, with  $r_0 = r_d = 1$ . For each input coordinate  $i = 1, \dots, d$ , the set of univariate functions  $\mathcal{F}_i = \{f_{(\alpha_{i-1}, \alpha_i)}^{(i)} : [a_i, b_i] \rightarrow \mathbb{R}, \alpha_{i-1} \in \{1, \dots, r_{i-1}\}, \alpha_i \in \{1, \dots, r_i\}\}$  are called *cores*. Each set can be viewed as a matrix-valued function and visualized as an array of univariate functions:

$$\mathcal{F}_i(x_i) = \begin{bmatrix} f_{1,1}^{(i)}(x_i) & \dots & f_{1,r_i}^{(i)}(x_i) \\ \vdots & \ddots & \vdots \\ f_{r_{i-1},1}^{(i)}(x_i) & \dots & f_{r_{i-1},r_i}^{(i)}(x_i) \end{bmatrix}.$$

Thus the evaluation of a function in FT format may equivalently be posed as a sequence of  $d-1$  vector-matrix products:

$$f(x_1, \dots, x_d) = \mathcal{F}_1(x_1) \dots \mathcal{F}_d(x_d). \quad (19)$$

The tensor-train decomposition differs from the canonical tensor decomposition by allowing a greater variety of interactions between neighboring dimensions through products of univariate functions in neighboring cores  $\mathcal{F}_i, \mathcal{F}_{i+1}$ . Furthermore, each of the cores contains  $r_{i-1} \times r_i$  univariate functions instead of a fixed number  $R$  for each  $\mathcal{F}_i^{CP}$  within the CP format.

The ranks of the FT decomposition of a function  $f$  can be bounded by the singular value decomposition (SVD) ranks of certain separated representations of  $f$ . Let  $\mathcal{X}_{\leq k} := [a_1, b_1] \times [a_2, b_2] \times \dots \times [a_k, b_k]$  and  $\mathcal{X}_{> k} := [a_{k+1}, b_{k+1}] \times \dots \times [a_d, b_d]$ , such that  $\mathcal{X} = \mathcal{X}_{\leq k} \times \mathcal{X}_{> k}$ . We then let  $f^i$  denote the  $i$ -separated representation of the function  $f$ , also called the  $i$ th *unfolding* of  $f$ :

$$f^i : \mathcal{X}_{\leq i} \times \mathcal{X}_{> i} \rightarrow \mathbb{R}, \text{ where}$$

$$f^i(\{x_1, \dots, x_i\}, \{x_{i+1}, \dots, x_d\}) = f(x_1, \dots, x_d).$$

The FT ranks of  $f$  are related to the SVD ranks of  $f^i$  via the following result.

**Theorem 4** (Ranks of approximately low-rank unfoldings, from Theorem 4.2 in [24]). *Suppose that the unfoldings  $f^i$  of the function  $f$  satisfy<sup>2</sup>, for all  $i = 1, \dots, d-1$ ,*

$$f^i = g^i + e^i, \quad \text{rank } g^i = r_i, \quad \|e^i\|_{L^2} = \varepsilon_i.$$

*Then a rank  $\mathbf{r} = [r_0, r_1, \dots, r_d]$  approximation  $\hat{f}$  of  $f$  in FT format may be obtained with bounded error<sup>3</sup>*

$$\int (f - \hat{f})^2 dx \leq \sum_{i=1}^{d-1} \varepsilon_i^2.$$

The ranks of an FT approximation of  $f$  can also be related to the Sobolev regularity of  $f$ , with smoother functions having

<sup>2</sup>The rank condition on  $g^i$  is based on the *functional* SVD, or Schmidt decomposition, a continuous analogue of the discrete SVD.

<sup>3</sup>In this paper, integrals  $\int f dx$  will always be with respect to the Lebesgue measure. For example,  $\int f(x_i) dx_i$  should be understood as  $\int f(x_i) \mu(dx_i)$ , and similarly  $\int f(x) dx = \int f(x) \mu(dx)$ .

faster approximation rates in FT format; precise results are given in [8]. These regularity conditions are sufficient but not necessary, however; according to Theorem 4 even discontinuous functions can have low FT ranks, if their associated unfoldings exhibit low rank structure.

## B. Cross approximation and rounding

The proof of Theorem 4 is constructive and results in an algorithm that allows one to decompose a function into its FT representation using a sequence of SVDs. While this algorithm, referred to as TT-SVD [47], exhibits certain optimality properties, it encounters the curse of dimensionality associated with computing the SVD of each unfolding  $f^i$ .

To remedy this problem, Oseledets [49] replaces the SVD with a cross approximation algorithm for computing CUR/skeleton decompositions [20], [43], [61] of matrices, within the overall context of compressing a tensor into TT format. Similarly, our previous work [24] employs a continuous version of cross approximation to compress a function into FT format. If each unfolding function  $f^i$  has a finite SVD rank, then these cross approximation algorithms can yield exact reconstructions.

The resulting algorithm only requires the evaluation of univariate function fibers. Fibers are multidimensional analogues of matrix rows and columns, and they are obtained by fixing all dimensions except one [35]. Each FT core can be viewed as a collection of fibers of the corresponding dimension, and the cross approximation algorithm only requires  $\mathcal{O}(nr^2)$  evaluations of  $f$ , where  $n$  represents the number of parameters used to represent each univariate function in each core.

While low computational costs make this algorithm attractive, there are two downsides. First, there are no convergence guarantees for the cross approximation algorithm when the unfolding functions are not of finite rank. Second, even for finite-rank tensors, the algorithm requires the specification of upper bounds on the ranks. If the upper bounds are set too low, then errors occur in the approximation; if the upper bounds are set too high, however, then too many function evaluations are required.

We mitigate the second downside, the specification of ranks, using a rank adaptation scheme. The simplest adaptation scheme, and the one we use for the experiments in this paper, is based on TT-rounding. The idea of TT-rounding [49] is to approximate a tensor in TT format  $\mathcal{T}$  by another tensor in TT format  $\hat{\mathcal{T}}$  to a tolerance  $\epsilon$ , i.e.,

$$\|\mathcal{T} - \hat{\mathcal{T}}\| \leq \epsilon \|\mathcal{T}\|.$$

The benefit of such an approximation is that a tensor in TT format can often be well *approximated* by another tensor with lower ranks if one allows for a relative error  $\epsilon$ . Furthermore, performing the rounding operation requires  $\mathcal{O}(nr^3)$  operations, where  $r$  refers to the rank of  $\mathcal{T}$ .

In our continuous context, we use a continuous analogue of TT-rounding [24] to aid in rank adaptation. The rank adaptation algorithm we use requires the following steps:

- 1) Estimate an upper bound on each rank  $r_i$

- 2) Use cross approximation to obtain a corresponding FT approximation  $\hat{f}$
- 3) Perform FT-rounding [24] with tolerance  $\epsilon$  to obtain a new  $\hat{f}$
- 4) If ranks of  $\hat{f}$  are not smaller than the ranks of  $\hat{f}$ , increase the ranks and go to step 2.

If all ranks are not rounded down, we may have underspecified the proper ranks in Step 1. In that case, we increase the upper bound estimate and retry.

The pseudocode for this algorithm is provided by Algorithm 2. Within this algorithm, there are calls to cross approximation (`cross-approx`) and to rounding (`ft-round`). A detailed description of these algorithms can be found in [24]. Algorithm 2 requires four inputs and produces one output. The inputs are: a black-box function that is to be approximated, a cross approximation tolerance  $\delta_{\text{cross}}$ , a rounding tolerance  $\epsilon_{\text{round}}$ , and an initial rank estimate. The output of `ft-rankadapt` is a separable approximation in FT format. We abuse notation by defining  $\epsilon = (\delta_{\text{cross}}, \epsilon_{\text{round}})$ , and refer to this procedure as  $\hat{f} = \text{ft-rankadapt}(f, \epsilon)$ .

---

**Algorithm 2** `ft-rankadapt`: Function-train approximation with rank adaptation [24]

---

**Require:** A  $d$ -dimensional black-box function  $f : \mathcal{X} \rightarrow \mathbb{R}$ ;  
Cross-approximation tolerance  $\delta_{\text{cross}}$ ;  
Size of rank increase `kickrank`;  
Rounding accuracy  $\epsilon_{\text{round}}$ ;  
Initial rank estimates  $\mathbf{r}$

**Ensure:** Approximation  $\hat{f}$  such that a rank increase and rounding does not change ranks.

```

1:  $\hat{f} = \text{cross-approx}(f, \mathbf{r}, \delta_{\text{cross}})$ 
2:  $\hat{f}_r = \text{ft-round}(\hat{f}, \epsilon_{\text{round}})$ 
3:  $\hat{\mathbf{r}} = \text{rank}(\hat{f}_r)$ 
4: while  $\exists i$  s.t.  $\hat{r}_i = r_i$  do
5:   for  $k = 1 \rightarrow d - 1$  do
6:      $\mathbf{r}_k = \hat{\mathbf{r}}_k + \text{kickrank}$ 
7:      $\hat{f} = \text{cross-approx}(f, \mathbf{r}, \delta_{\text{cross}})$ 
8:   end for
9:    $\hat{f}_r = \text{ft-round}(\hat{f}, \epsilon_{\text{round}})$ 
10:   $\hat{\mathbf{r}} = \text{rank}(\hat{f}_r)$ 
11: end while
12:  $\hat{f} = \hat{f}_r$ 

```

---

This algorithm requires  $\mathcal{O}(nr^2)$  evaluations of the function  $f$ , and it requires  $\mathcal{O}(nr^3)$  operations in total.

## V. LOW-RANK COMPRESSED DYNAMIC PROGRAMMING

In this section, we propose novel dynamic programming algorithms based on the compressed continuous computation framework described in the previous section. Specifically, we describe how to represent value functions in FT format in Section V-A, and how to perform FT-based versions of the value iteration, policy iteration, and multigrid algorithms in Sections V-B, V-C, and V-D, respectively.

### A. FT representation of value functions

The Markov chain approximation method approximates a continuous-space stochastic control problem with a discrete-

state Markov decision process. In this framework, the continuous value functions  $w$  are approximated by their discrete counterparts  $w^h$ . In order to leverage low-rank decompositions, We focus our attention on situations where the discrete value functions represent the cost of discrete MDPs defined through a tensor-product discretization of the state space, and therefore,  $w^h$  can be interpreted as a  $d$ -way array. In order to combat the curse of dimensionality associated with storing and computing  $w^h$ , we propose using the function-train decomposition.

Let  $\mathcal{X} = \mathcal{X}_1 \times \mathcal{X}_2 \times \dots \times \mathcal{X}_d$  denote a tensor-product of intervals as described in Section IV-A. Recall that the space of the  $i$ th state variable is  $\mathcal{X}_i = [a_i, b_i]$ , for  $a_i, b_i \in \mathbb{R}$  that have the property  $a_i < b_i$ . A *tensor-product discretization*  $\mathcal{X}^h$  of  $\mathcal{X}$  involves discretizing each dimension into  $n$  nodes to form  $\mathcal{X}_i^h \subset \mathcal{X}_i$  where

$$\mathcal{X}_i^h = \{z_1^{(i)}, z_2^{(i)}, \dots, z_n^{(i)}\}, \quad \text{where } z_k^{(i)} \in \mathcal{X}_i \text{ for all } k.$$

A discretized value function  $w^h$  can therefore be viewed as a vector with  $n^d$  elements.

The Markov chain approximation method guarantees that the solution to the discrete MDP approximates the solution to the original stochastic optimal control problem, i.e.,  $v^h \approx v$ , for small enough discretizations. This approximation, however, is ill-defined since  $v^h$  is a multidimensional array and  $v$  is a multivariate function. Furthermore, a continuous control law requires the ability to determine the optimal control for a system when it is in a state that is not included in the discretization. Therefore, it will be necessary to use the discrete value function  $w^h \in \mathcal{X}^h$  to develop a value function that is a mapping from the continuous space  $\mathcal{X}$  to the reals.

We will slightly abuse notation and interpret  $w^h$  as both a value function of the discrete MDP and as an *approximation* to value function of the continuous system. Furthermore, when generating this continuous-space approximation, we are restricted to evaluations located only within the tensor-product discretization. In this sense, we can think of  $w^h$  both as an array  $w^h : \mathcal{X}^h \rightarrow \mathbb{R}$  in the sense that it has elements

$$w^h[i_1, \dots, i_d] \approx w(z_1^{(i_1)}, \dots, z_d^{(i_d)}),$$

and simultaneously as a function  $w^h : \mathcal{X} \rightarrow \mathbb{R}$  where

$$w^h(z_1^{(i_1)}, \dots, z_d^{(i_d)}) \approx w(z_1^{(i_1)}, \dots, z_d^{(i_d)}).$$

Now, recall that the FT representation of a function  $f$  is given by Equation (18) and defined by the set of FT cores  $\{\mathcal{F}_i\}$  for  $i = 1, \dots, d$ . Each of these cores is a matrix-valued function  $\mathcal{F}_i : \mathcal{X}_i \rightarrow \mathbb{R}^{r_{i-1} \times r_i}$  that can be represented as a two-dimensional array of scalar-valued univariate functions:

$$\mathcal{F}_i(x_i) = \begin{bmatrix} f_{1,1}^{(i)}(x_i) & \dots & f_{1,r_i}^{(i)}(x_i) \\ \vdots & \ddots & \vdots \\ f_{r_{i-1},1}^{(i)}(x_i) & \dots & f_{r_{i-1},r_i}^{(i)}(x_i) \end{bmatrix}.$$

Using this matrix-valued function representation for the cores, the evaluation of a function in the FT format can be expressed as  $f(x_1, \dots, x_d) = \mathcal{F}_1(x_1) \dots \mathcal{F}_d(x_d)$ .



Since we can effectively compute  $w^h$  through evaluations at uniformly discretized tensor-product grids, we employ a *nodal* representation of each scalar-valued univariate function

$$f_{\alpha_1, \alpha_2}^{(i, h)}(x_i) = \sum_{\ell=1}^n a_{\alpha_1, \alpha_2, \ell}^{(i, h)} \phi_{\ell}^h(x_i), \quad (20)$$

where  $a_{\alpha_1, \alpha_2, \ell}^{(i, h)}$  are the coefficients of the expansion, and the basis functions  $\phi_{\ell}^h : \mathcal{X}_i \rightarrow \mathbb{R}$  are hat functions:

$$\phi_{\ell, h}(x_i) = \begin{cases} 0 & \text{if } x_i < z_{\ell-1}^{(i)} \text{ or } x_i > z_{\ell+1}^{(i)} \\ 1 & \text{if } x_i = z_{\ell}^{(i)} \\ \frac{x_i - z_{\ell-1}^{(i)}}{z_{\ell}^{(i)} - z_{\ell-1}^{(i)}} & \text{if } z_{\ell-1}^{(i)} \leq x_i < z_{\ell}^{(i)} \\ \frac{z_{\ell+1}^{(i)} - x_i}{z_{\ell+1}^{(i)} - z_{\ell}^{(i)}} & \text{if } z_{\ell}^{(i)} < x_i \leq z_{\ell+1}^{(i)} \end{cases}.$$

Note that these basis functions yield a *linear element* interpolation<sup>4</sup> of the function when evaluating it for a state not contained within  $\mathcal{X}^h$ . We will denote the FT cores of the value functions for discretization level  $h$  as  $\mathcal{F}_i^h$ . Finally, evaluating  $w^h$  anywhere within  $\mathcal{X}$  requires evaluating a sequence of matrix-vector products

$$w^h(x_1, x_2, \dots, x_d) = \mathcal{F}_1^h(x_1) \mathcal{F}_2^h(x_2) \dots \mathcal{F}_d^h(x_d),$$

for all  $x_i \in \mathcal{X}_i$ .

### B. FT-based value iteration algorithm

In prior work [23], we introduced a version of VI in which each update given by Equation (10) was performed by using a low-rank tensor interpolation scheme to selectively choose a small number of states  $z \in \mathcal{X}^h$ . Here we follow a similar strategy, except we use the continuous space approximation algorithm described in Algorithm 2, and denoted by `ft-rankadapt`, to accomodate our continuous space approximation. By seeking a low-rank representation of the value function, we are able to avoid visiting every state in  $\mathcal{X}^h$  and achieve significant computational savings. The pseudocode for low-rank VI is provided by Algorithm 3. In this algorithm,  $v_k^h$  denotes the value function approximation during the  $k$ -th iteration.

---

#### Algorithm 3 FT-based Value Iteration (FTVI)

---

**Require:** Termination criterion  $\delta_{\max}$ ;

`ft-rankadapt` tolerances  $\epsilon = (\delta_{\text{cross}}, \epsilon_{\text{round}})$ ;

Initial cost function in FT format  $v_0^h$ ;

Convergence tolerance  $\delta_{\max}$

**Ensure:** Residual  $\delta = \|v_k^h - v_{k-1}^h\|^2 < \delta_{\max}$

---

1:  $k = 0$

2: **repeat**

3:  $v_{k+1}^h = \text{ft-rankadapt}(T^h(v_k^h), \epsilon)$

4:  $k \leftarrow k + 1$

5:  $\delta = \|v_k^h - v_{k-1}^h\|^2$

6: **until**  $\delta < \delta_{\max}$

---

<sup>4</sup>If we would have chosen a piecewise constant reconstruction, then for all intents and purposes the FT would be equivalent to the TT. Indeed we have previously performed such an approximation for the value function [23].

The update step 3 treats the VI update as a black box function into which one feeds a state and obtains an updated cost. After  $k$  steps of FT-based VI we can obtain a policy as the minimizer of  $T^h(v_k^h)(z)$ , for any state  $z \in \mathcal{X}$ .

### C. FT-based policy iteration algorithm

As part of the policy iteration algorithm, for each  $\mu_k$ , one needs to solve Equation (12) for the corresponding value function  $w_k^h$ . This system of equations has an equivalent number of unknowns as states in  $\mathcal{X}^h$ . Hence, the number of unknowns grows exponentially with dimension, for tensor-product discretizations. In order to efficiently solve this system in high dimensions we seek *low-rank* solutions. A wide variety of low-rank linear system solvers have recently been developed that can potentially be leveraged for this task [16], [48].

We focus on optimistic policy iteration, where we utilize the contractive property of  $T_{\mu_k}^h$  to solve Equation (12) approximately, using  $n_{fp}$  fixed point iterations. See Section III-C4. We leverage the low-rank nature of each intermediate value  $w_k^h$  by interpolating a new value function for each of these iterations. Notice that this iteration in Equation (9) is much less expensive than the value iteration in Equation (10), because it does not involve any minimization.

The pseudocode for the FT-based optimistic policy iteration is provided in Algorithm 4. In Line 3, we represent a policy  $\mu_k$  *implicitly* through a value function. We make this choice, instead of developing a low-rank representation of  $\mu_k$ , because the policies are generally not low-rank, in our experience! Indeed, discontinuities can arise due to regions of uncontrollability, and these discontinuities increase the rank of the policy. Instead, to evaluate an implicit policy  $\mu_k$  at a location  $z$ , one is required to solve the optimization problem in Equation (11) using the fixed value function  $w_{k-1}^h$ . However, one can store the policy evaluated at nodes visited by the cross approximation algorithm to avoid recomputing them during each iteration of the loop in Line 5. Since the number of nodes visited during the approximation stage should be relatively low, this does not pose an excessive algorithmic burden.

---

#### Algorithm 4 FT-based Optimistic Policy Iteration (FTPI)

---

**Require:** Termination criterion  $\delta_{\max}$ ;

`ft-rankadapt` tolerances  $\epsilon = (\delta_{\text{cross}}, \epsilon_{\text{round}})$ ;

Initial value function in FT format  $w_0^h$ ;

Number of FP sub-iterations  $n_{fp}$

1:  $k = 1$

2: **repeat**

3:  $\mu_k = \text{ImplicitPolicy}\left(\arg \min_{\mu} [T_{\mu}^h(w_{k-1}^h)]\right)$

4:  $w_k^h = w_{k-1}^h$

5: **for**  $\ell = 1$  **to**  $n_{fp}$  **do**

6:  $w_k^h = \text{ft-rankadapt}(T_{\mu}^h(w_k^h), \epsilon)$

7: **end for**

8:  $\delta = \|w_k^h - w_{k-1}^h\|^2$

9:  $k \leftarrow k + 1$

10: **until**  $\delta < \delta_{\max}$

---

#### D. FT-based multigrid algorithm

In this section, we propose a novel FT-based multigrid algorithm. Recall that the first ingredient of multigrid is a prolongation operator  $I_h^{2h}$  which takes functions defined on the grid  $\mathcal{X}^h$  into a coarser grid  $\mathcal{X}^{2h}$ . The second ingredient is an interpolation operator  $I_{2h}^h$  which interpolates functions defined on  $\mathcal{X}^{2h}$  onto the functions defined on  $\mathcal{X}^h$ .

Many of the common operators used for  $I_h^{2h}$  and  $I_{2h}^h$  can take advantage of the low rank structure of any functions on which they are operating. In particular, performing these operations on function in low-rank format often simply requires performing their one-dimensional variants onto each univariate scalar-valued function of its FT core. Let us describe the FT-based prolongation and interpolation operators.

The *prolongation operator* that we use picks out values common to both  $\mathcal{X}^h$  and  $\mathcal{X}^{2h}$  according to

$$(I_h^{2h} w^h)(z) = w^{2h}(z), \quad \forall z \in \mathcal{X}^{2h}$$

This operator requires a constant number of computational operations, only access to memory. Furthermore, the coefficients of  $f_{\alpha_1, \alpha_2}^{(i, h)}$  are reused to form the coefficients of  $f_{\alpha_1, \alpha_2}^{(i, 2h)}$ . See Equation (20). Suppose that fine grid is  $\mathcal{X}_i^h = \{z_1^{(i)}, z_2^{(i)}, \dots, z_n^{(i)}\}$  and the coarse grid  $\mathcal{X}_i^{2h} = \{z_{2l-1}^{(i)}\}_{l=1, \dots, n/2}$  consists of half the nodes ( $|\mathcal{X}_i^{2h}| = n/2$ ). Then the univariate functions making up the cores of  $w^{2h}$  are

$$\begin{aligned} f_{j,k}^{(i, 2h)}(x_i) &= \sum_{\ell=1}^{n/2} a_{j,k,\ell}^{(i, 2h)} \phi_\ell^{2h}(x_i), \\ &= \sum_{\ell=1}^{n/2} a_{j,k,2\ell-1}^{(i, h)} \phi_\ell^{2h}(x_i), \end{aligned}$$

where  $\phi_\ell^{2h}$  are the functions defined on the coarser grid. In the second line, we use every other coefficient from the finer grid as coefficients of the corresponding coarser grid function.<sup>5</sup>

The *interpolation operator* arises from the interpolation that the FT performs, and in our case, the use of hat functions leads to a linear interpolation scheme. This means that if the scalar-valued univariate functions making up the cores of  $w^{2h}$  are represented in a nodal basis obtained at the tensor product grid  $\mathcal{X}^{2h} = \mathcal{X}_1^{2h} \times \dots \times \mathcal{X}_d^{2h}$ , then we obtain a nodal basis with twice the resolution defined on  $\mathcal{X}^h = \mathcal{X}_1^h \times \dots \times \mathcal{X}_d^h$  through interpolation of each core. This operation requires interpolating of each of the univariate functions making up the cores of  $w^{2h}$  onto the fine grid. Thus  $w^h$  becomes an FT with cores consisting of the univariate functions

$$\begin{aligned} f_{j,k}^{(i, h)}(x_i) &= \sum_{\ell=1}^n a_{j,k,\ell}^{(i, h)} \phi_\ell^h(x_i), \\ &= \sum_{\ell=1}^n f_{j,k}^{(i, 2h)}(z_\ell^{(i)}) \phi_\ell^h(x_i), \end{aligned}$$

<sup>5</sup>Alternatively, one can choose more regular nodal basis functions, such as splines. For other basis functions, there may be more natural prolongation and interpolation operators. We choose hat functions in this paper, because we observe that they are well behaved in the face of discontinuities or extreme nonlinearities often encountered in the solution of the HJB equation.

where we note that in the second equation we use evaluations of the coarser basis functions to obtain the coefficients of the new basis functions. In summary, both operators can be applied to each FT core of the value function separately. Both of these operations, therefore, scale linearly with dimension.

The FT-based two-level V-grid algorithm is provided in Algorithm 5. In this algorithm, an approximate solution for each equation is obtained by  $\ell_i$  fixed point iterations at grid level  $h_i$ . An extension to other grid cycles and multiple levels of grids is straightforward and can be performed with all of the same operations.

---

#### Algorithm 5 Two-level FT-based V-grid

---

**Require:** Discretization levels  $\{h_1, h_2\}$  such that  $h_1 < h_2$ ;  
Number of iterations at each level  $\{\ell_1, \ell_2\}$ ;  
Initial value function  $\hat{w}_0^h$ ;  
Policy  $\mu$ ;  
ft-rankadapt tolerances  $\epsilon = (\delta_{\text{cross}}, \epsilon_{\text{round}})$

- 1:  $k = 0$
- 2: **repeat**
- 3:   **for**  $\ell = 1, \dots, \ell_1$  **do**
- 4:      $\hat{w}_k^{h_1} \leftarrow \text{ft-rankadapt}(T_\mu^{h_1}(w_k^{h_1}), \epsilon)$
- 5:   **end for**
- 6:    $r^h = \text{ft-rankadapt}(T_{c,\mu}^h(\hat{w}_k^{h_1}), \epsilon) - \hat{w}_k^{h_1}$
- 7:    $\Delta w^{h_2} = 0$
- 8:   **for**  $\ell = 1, \dots, \ell_2$  **do**
- 9:      $\Delta w_{n+1}^{h_2} \leftarrow \text{ft-rankadapt}(T_\mu^{h_2}(\Delta w_n^{h_2}; I_{h_1}^{h_2} r^h), \epsilon)$   
      # See (16)
- 10:   **end for**
- 11:    $\hat{w}_{k+1}^{h_1} = \hat{w}_k^{h_1} + I_{h_2}^{h_1} \Delta w^{h_2}$
- 12:    $k \leftarrow k + 1$
- 13: **until** convergence

---

## VI. ANALYSIS

In this section, we analyze the convergence properties and the computational complexity of the FT-based algorithms proposed in the previous section. First, we prove the convergence and accuracy of approximate fixed-point iteration methods in Section VI-A. Then, in Section VI-B, we apply these result to prove the convergence of the proposed FT-based algorithms we discuss computational complexity in Section VI-C.

The algorithms discussed in the previous section all rely on performing cross approximation of a function with a relative accuracy tolerance  $\epsilon$ . Recall, from Section IV-B that cross approximation yields an exact reconstruction only when the value function has finite rank unfoldings. In this case, rounding to a relative error  $\epsilon$  also guarantees that we achieve an  $\epsilon$ -accurate solution. In what follows, we provide conditions under which FT-based dynamic programming algorithms converge, assuming the rank-adaptive cross approximation algorithm ft-rankadapt provides an approximation with at most  $\epsilon$  error.

#### A. Convergence of approximate fixed-point iterations

We start by showing that a small relative error made during each step of the relevant fixed-point iterations result in a bounded overall approximation error.

We begin recalling the contraction mapping theorem.

**Theorem 5** (Contraction mapping theorem; Proposition B.1 by Bertsekas [6]). *Let  $\mathcal{R}$  be a complete vector space and  $\mathcal{A}$  be a closed subset. Then if  $f : \mathcal{R} \rightarrow \mathcal{R}$  is a contraction mapping with modulus  $\gamma \in (0, 1)$ , there exists a unique  $w \in \mathcal{A}$  such that  $w = f(w)$ . Furthermore, the sequence defined by  $w_0 \in \mathcal{A}$  and the iteration  $w_k = f(w_{k-1})$  converges to  $w$  for any  $w \in \mathcal{A}$  according to*

$$\|w_k - w\| \leq \gamma^k \|w_0 - w\|, \quad k = 1, 2, \dots$$

Theorem 5 can be used, for example, to prove the convergence of the value iteration algorithm, when the operator  $T^h$  is a contraction mapping. The proposed FT-based algorithms are based on *approximations* to contraction mappings. To prove their convergence properties, it is important to understand when approximate fixed-point iterations converge and the accuracy which they attain. Lemma 1 below addresses the convergence of approximate fixed-point iterations.

**Lemma 1** (Convergence of approximate fixed-point iterations). *Let  $\mathcal{R}^h$  be a closed subset of a complete vector space. Let  $f : \mathcal{R}^h \rightarrow \mathcal{R}^h$  be a contractive mapping with modulus  $\gamma \in (0, 1)$  and fixed point  $w^h$ . Let  $\tilde{f} : \mathcal{R}^h \rightarrow \mathcal{R}^h$  be an approximate mapping such that*

$$\|\tilde{f}(w') - f(w')\| \leq \epsilon \|f(w')\|, \quad \forall w' \in \mathcal{R}^h, \quad (21)$$

for  $\epsilon > 0$ . Then, the sequence defined by  $w_0^h \in \mathcal{R}^h$  and the iteration  $w_k^h = \tilde{f}(w_{k-1}^h)$  for  $k = 1, 2, \dots$  satisfies

$$\|w_k^h - w^h\| \leq \epsilon \frac{1 - (\gamma\epsilon + \gamma)^k}{1 - (\gamma\epsilon + \gamma)} \|w^h\| + (\gamma\epsilon + \gamma)^k \|w_0^h - w^h\|.$$

*Proof.* The proof is a standard contraction argument. We begin by bounding the difference between the  $k$ -th iterate and the fixed point  $w^h$ :

$$\begin{aligned} \|w_k^h - w^h\| &= \|w_k^h - f(w_{k-1}^h) + f(w_{k-1}^h) - w^h\| \\ &\leq \|w_k^h - f(w_{k-1}^h)\| + \|f(w_{k-1}^h) - w^h\| \\ &\leq \epsilon \|f(w_{k-1}^h)\| + \gamma \|\widehat{w}_{k-1}^h - w^h\| \\ &\leq \epsilon \|f(w_{k-1}^h) - w^h\| + \epsilon \|w^h\| + \gamma \|w_{k-1}^h - w^h\| \\ &\leq (\gamma\epsilon + \gamma) \|w_{k-1}^h - w^h\| + \epsilon \|w^h\|, \end{aligned}$$

where the second inequality comes from the triangle inequality, the third comes from Equation (21) and contraction, the fourth inequality arises again from the triangle inequality, the final inequality arises from contraction. Using recursion results in

$$\begin{aligned} \|w_k^h - w^h\| &\leq (\gamma\epsilon + \gamma) \left[ (\gamma\epsilon + \gamma) [(\gamma\epsilon + \gamma) \dots + \epsilon \|w^h\|] \right. \\ &\quad \left. + \epsilon \|w^h\| \right] + \epsilon \|w^h\| \\ &= \epsilon \|w^h\| \left[ \sum_{\ell=0}^{k-1} (\gamma\epsilon + \gamma)^\ell \right] + \\ &\quad (\gamma\epsilon + \gamma)^k \|\widehat{w}_0^h - w^h\|. \end{aligned}$$

Evaluating the sum of a geometric series,

$$\|\widehat{w}_k^h - w^h\| \leq \epsilon \frac{1 - (\gamma\epsilon + \gamma)^k}{1 - (\gamma\epsilon + \gamma)} \|w^h\| + (\gamma\epsilon + \gamma)^k \|\widehat{w}_0^h - w^h\|,$$

we reach the desired result.  $\square$

Notice that, as expected, when  $\epsilon = 0$  and  $k \rightarrow \infty$  this result yields that the iterates  $w_k^h$  converge to the fixed point  $w^h$ . Second, the condition  $\gamma\epsilon + \gamma < 1$  is required to avoid divergence. This condition effectively states that, if the contraction modulus is small enough, a larger approximation error may be incurred. On the other hand, if the contraction modulus is large, then the approximation error must be small. In other words, this requirement can be thought as a condition for which the approximation can remain a contraction mapping; larger errors can be tolerated when the original contraction mapping has a small modulus, and vice-versa.

The following result is an alternative to the previous lemma. It is more flexible in the size of error that it allows, so long as each iterate is bounded.

**Lemma 2** (Convergence of approximate fixed-point iterations with a boundedness assumption). *Let  $\mathcal{R}^h$  be a closed subset of a complete vector space. Let  $f : \mathcal{R}^h \rightarrow \mathcal{R}^h$  be a contractive mapping with modulus  $\gamma \in (0, 1)$  and fixed point  $w^h$ . Let  $\tilde{f} : \mathcal{R}^h \rightarrow \mathcal{R}^h$  be an approximate mapping such that*

$$\|\tilde{f}(w') - f(w')\| \leq \epsilon \|f(w')\|, \quad \forall w' \in \mathcal{R}^h, \quad (22)$$

for  $\epsilon > 0$ . Let  $w_0^h \in \mathcal{R}^h$ . Define a sequence the sequence  $\{w_k^h\}$  according to the iteration  $w_k^h = \tilde{f}(w_{k-1}^h)$  for  $k = 1, 2, \dots$ . Assume that  $\|w_k^h\| \leq \rho_1 < \infty$  so that  $\|f(w_k^h)\| \leq \rho < \infty$ . Then  $\{w_k^h\}$  satisfies

$$\|w_k^h - f^{[k]}(w_0^h)\| \leq \frac{1 - \gamma^k}{1 - \gamma} \epsilon \rho, \quad \forall k, \quad (23)$$

so that,

$$\lim_{k \rightarrow \infty} \|w_k^h - w^h\| \leq \frac{\epsilon \rho}{1 - \gamma}, \quad (24)$$

where  $f^{[k]}$  denotes  $k$  applications of the mapping  $f$ .

*Proof.* The strategy for this proof again relies on standard contraction and triangle inequality arguments. Furthermore, it follows closely the proof of Proposition 2.3.2 (Error Bounds for Approximate VI) work by Bertsekas [6]. In that work, the proposition provides error bounds for approximate VI when an absolute error (rather than a relative error) is made during each approximate FP iteration. The assumption of boundedness that we use here will allow us to use the same argument as the one made by Bertsekas.

Equation (22) implies

$$\|w_k^h - f(w_{k-1}^h)\| = \|\tilde{f}(w_{k-1}^h) - f(w_{k-1}^h)\| \leq \epsilon \|f(w_{k-1}^h)\|. \quad (25)$$

Recall  $f^{[k]}(w')$  denotes  $k$  applications of the operator  $f$ , i.e.,

$$\begin{aligned} f^{[k]}(w') &= f^{[k-1]}(f(w')) = f^{[k-2]}(f(f(w'))) = \dots \\ &= f(f(f(\dots f(w')))). \end{aligned}$$

Using the triangle inequality, contraction, and Equation (25),

$$\begin{aligned} \|w_k^h - f^{[k]}(w_0^h)\| &\leq \|w_k^h - f(w_{k-1}^h)\| \\ &\quad + \|f(w_{k-1}^h) - f^{[2]}(w_{k-2}^h)\| + \dots \\ &\quad + \|f^{[k-1]}(w_1^h) - f^{[k]}(w_0^h)\| \\ &\leq \epsilon \|f(w_{k-1}^h)\| + \gamma \epsilon \|f(w_{k-2}^h)\| + \dots \\ &\quad + \gamma^{k-1} \epsilon \|f(w_0^h)\|. \end{aligned}$$

The boundedness assumption yields

$$\|w_k^h - f^{[k]}(w_0^h)\| \leq \epsilon \rho + \gamma \epsilon \rho + \dots + \gamma^{k-1} \epsilon \rho.$$

We then evaluate the sum of a geometric series,

$$\|w_k^h - f^{[k]}(w_0^h)\| \leq \frac{1 - \gamma^k}{1 - \gamma} \epsilon \rho.$$

Taking the limit  $k \rightarrow \infty$  and using Theorem 5, where  $\lim_{k \rightarrow \infty} f^{[k]}(w_0^h) = w^h$ , yields Equation (24).  $\square$

Equation (23) shows the difference between iterates of exact fixed-point iteration and approximate fixed-point iteration, and indicates that this difference can not grow larger than  $\epsilon \rho$ . Furthermore, this result does not require any assumption on the relative error  $\epsilon$ .

In the next section, we use these intermediate results to prove that the proposed FT-based dynamic programming algorithms converge under certain conditions.

### B. Convergence of FT-based fixed-point iterations

In order for the above theorems to be applicable to the case of low-rank approximation using the `ft-rankadapt` algorithm, we need to be able to explicitly bound the error committed during each iteration of FT-based value iteration, and each subiteration within the optimistic policy iteration algorithm. In order to strictly adhere to the intermediate results of the previous section, we focus our attention to functions with finite FT rank. This assumption is formalized below.

**Assumption 4** (Bounded ranks of FT-based FP iteration). *Let  $\mathcal{R}^h$  be a closed subset of a complete vector space. Let  $f : \mathcal{R}^h \rightarrow \mathcal{R}^h$  be a contractive mapping with modulus  $\gamma \in (0, 1)$  and fixed point  $w^h$ . The sequence defined by the initial condition  $w_0^h \in \mathcal{R}^h$  and the iteration  $w_k^h = \text{ft-rankadapt}(f(w_{k-1}^h), \epsilon)$  has the property that the functions  $\{f(w_k^h)\}$  have finite FT rank. In other words, each of the unfoldings of  $f(w_k^h)$  have rank bounded by some  $r < \infty$ ,*

$$\text{rank}[f(w_k^h)(\{x_1, \dots, x_l\}; \{x_{l+1} \dots x_d\})] < r < \infty,$$

for  $l = 1, \dots, d - 1$ .

Since we are approximating a function with finite FT ranks at each step of the fixed-point iteration under Assumption 4, the `ft-rankadapt` algorithm converges. Thus, this assumption leads to the satisfaction of the conditions required by Lemmas 1 and 2. Then, the following result is immediate.

**Theorem 6** (Convergence of the FT-based value iteration algorithm). *Let  $\mathcal{R}^h$  be a closed subset of a complete vector space. Let the operator  $T^h$  of (8) be a contraction mapping with modulus  $\gamma$  and fixed point  $v^h$ . Define a sequence of*

*functions according to an initial function  $v_0^h \in \mathcal{R}^h$  and the iteration  $v_k^h = \text{ft-rankadapt}(T^h(v_{k-1}^h), \epsilon)$ . Assume Assumption 4. Furthermore, assume that  $\|w_k^h\| \leq \rho_1 < \infty$  so that  $\|f(w_k^h)\| \leq \rho < \infty$ . Then, FT-based VI converges according to*

$$\lim_{k \rightarrow \infty} \|v_k^h - v^h\| \leq \frac{\epsilon \rho}{1 - \gamma}.$$

### C. Complexity of FT-based dynamic programming algorithms

Suppose that each dimension is discretized into  $n$  nodes, then recall that Algorithm 2 consists of a single interpolation of black box function having all FT ranks equal to  $r$  that requires  $\mathcal{O}(nr^2)$  evaluations of the black box function for cross approximation and a rounding step that requires  $\mathcal{O}(nr^3)$  operations.

Suppose we use the upwind differencing scheme described in Section III-B1. Then, the complexity of one step of an approximate fixed-point iteration can be characterized as follows.

**Proposition 1** (Complexity of the evaluation of Equation (6)). *Let the evaluation of stage cost  $g^h$ , drift  $B$ , and diffusion  $\mathcal{D}$  require  $n_{\text{op}}$  operations. Let the discretization  $\mathcal{X}^h$  of the MCA method arise from a tensor product of  $n$ -node discretizations of each dimension of the state space. Let the resulting transition probabilities  $p^h(z, z' | \bar{u})$  be computed according to the upwind scheme described in Section III-B1. Furthermore, let the value function  $w_k^h$  have ranks  $\mathbf{r} = [r_0, r_1, \dots, r_d]$  where  $r_0 = r_d = 1$  and  $r_k = r$  for  $k = 1 \dots d - 1$ . Then, evaluating*

$$g^h(z, \bar{u}) + \gamma \sum_{z' \in \mathcal{X}^h} p^h(z, z' | \bar{u}) w_k^h(z')$$

*for a fixed  $z \in \mathcal{X}^h$  and  $\bar{u} \in \mathcal{U}$  requires  $\mathcal{O}(n_{\text{op}} + d^2 nr^2)$  operations.*

*Proof.* In the specified upwind discretization scheme, there exist  $2d$  neighbors to which the transition probabilities are non-zero. Furthermore, in Section III-B1, we showed that computing all of these transition probabilities requires  $\mathcal{O}(n_{\text{op}} + d)$  operations. The evaluation of the cost of each neighbor,  $w_k^h(z')$ , requires  $\mathcal{O}(dnr^2)$  evaluations. Since this evaluation is required at  $2d$  neighbors, a conservative estimate for this total cost is  $\mathcal{O}(d^2 nr^2)$ . Thus, the result is obtained by observing that the cost is dominated by the computation of the transition probabilities and the evaluation of value function at the neighboring grid points  $\square$

Using Proposition 1 and assuming that for each  $z$  the minimization over control  $\bar{u}$  requires  $\kappa$  evaluations of Equation (6), the following result is immediate.

**Theorem 7** (Computational complexity of the FT-based value iteration algorithm). *Let  $\mathcal{R}^h$  be a closed subset of a complete vector space. Let the operator  $T^h$  of Equation (8) be a contraction mapping with modulus  $\gamma$  and fixed point  $v^h$ . Let the evaluation of stage cost  $g^h$ , drift  $B$ , and diffusion  $\mathcal{D}$  corresponding to  $T^h$  require  $n_{\text{op}}$  operations. Let the discretization  $\mathcal{X}^h$  of the MCA method arise from a tensor product of  $n$ -node discretizations of each dimension of the state space.*

*Define a sequence of functions according to an initial function  $v_0^h \in \mathcal{R}^h$  and the iteration  $v_k^h =$*

$\text{ft-rankadapt}(T^h(v_{k-1}^h), \epsilon)$ . Assume that for each  $z \in \mathcal{X}_h$ , the minimization over control  $\bar{u} \in \mathcal{U}$  requires  $\kappa$  evaluations of Equation (6). Then, each iteration of FT-based VI involves cross approximation and rounding and requires

$$\mathcal{O}(dnr^2\kappa(n_{op} + d^2nr^2) + dnr^3)$$

operations.

We remark that the computational cost of the proposed FT-based value iteration algorithm grows polynomially with increasing dimensionality. Therefore, this algorithm mitigates the curse of dimensionality as long as the rank  $r$  of the problem does not grow exponentially with dimensionality.

## VII. COMPUTATIONAL RESULTS AND QUADCOPTER EXPERIMENTS

In this section, we demonstrate the proposed algorithm in a number of challenging examples. First, we demonstrate the algorithm on simple control problems involving linear dynamics, quadratic cost and Gaussian process noise in Section VII-A. While these problems are not high dimensional, we experiment with various parameters, such as the amount the noise, to see their effect on various problem variables, such as rank. In this manner, we gain insight into various problem dependent properties including the FT rank. Next, we consider four different dynamics with increasing dimensionality. In Section VII-B, we consider two models of car-like robots with dimension three and four. These models are nonlinear. In Section VII-C, we consider a widely-studied perching problem that features nonlinear underactuated dynamics with a seven-dimensional state space that is not affine in control. Finally, in Section VII-D, we consider a problem instance involving a quadcopter maneuvering through a tight window, using a nonlinear quadcopter model with a six-dimensional state space. We compute a full-state feedback controller, and demonstrate it on a quadcopter flying through a tight window using a motion-capture system for full state estimation.

The simulation results in this section are obtained using one core of a 32 GB Intel i7-5930K CPU clocked at 3.50GHz with a 64-bit Ubuntu 14.04 LTS operating system. We use the Compressed Continuous computation ( $C^3$ ) library [21] for FT-based compressed computation, and this library is BSD licensed and available through GitHub.

A multistart BFGS optimization algorithm, available within  $C^3$ , is used for generating a control for a particular state within simulation and real-time system operation in Section VII-E. In particular for any state  $z'$  encountered in our simulation/experiment, we use multistart BFGS to obtain the minimizer of Equation (4). Within the objective the compressed cost function is evaluated at the corresponding neighboring states, i.e.,  $v^h(z')$  is evaluated such that  $z'$  are neighbors determined by the MCA discretization of  $z$ .

The low-rank dynamic programming algorithms are provided in a stochastic control module that is released separately, also on GitHub [22].

### A. Linear-quadratic-Gaussian problems with bounded control

In this section, we investigate the effect of state boundary conditions and control bounds on a prototypical control system. The system has a bounded state space, linear dynamics and quadratic cost. However, it also has a bounded control space, thereby making analytical solutions difficult to obtain.

Consider the following stochastic dynamical system

$$\begin{aligned} dx_1 &= x_2 dt + \sigma_1 dw_1(t) \\ dx_2 &= u(t) dt + \sigma_2 dw_2(t). \end{aligned}$$

This stochastic differential equations represent a physical system with position  $x_1$  and velocity  $x_2$  with  $\mathcal{O} = (-2, 2)^2$ . The control input to this system is the acceleration  $u$ , and we consider different lower and upper bounds on the control space  $\mathcal{U} = [u_{lb}, u_{ub}]$ . We consider a discounted-cost infinite-horizon problem, with discount  $e^{-\frac{t}{10}}$ . The stage cost is

$$g(x, u) = x_1^2 + x_2^2 + u^2, \quad (26)$$

and the terminal cost is

$$\psi(x) = 100, \quad x \in \partial\mathcal{O} \quad (27)$$

We first solve the problem for various parameter values to gain insight to the problem. Next, we discuss numerical results summarizing the convergence of the algorithm.

1) *Parameter studies:* We present computational experiments that assess when low-rank cost functions arise and what factors affect ranks. Our first computational experiment studies the effects of the control boundary conditions on a problem with absorbing boundaries. This computational experiment is performed over several different  $u_{lb}, u_{ub}$  combinations and the resulting optimal value functions are shown in Figure 3. We utilized FT-based policy iteration. The Markov chain approximation was obtained using 60 points in each dimension.

Several phenomena are evident in Figure 3. When the range of the valid controls is wide, the value function is able to achieve smaller values, i.e., the blue region (indicating small costs) is larger with a wider control range. This characteristic is expected since the region from which the state can avoid the boundary is larger when more control can be exercised. However, the alignment of the value function, with the diagonal stretching from the upper left to the upper right, is the same for all of the test cases. Only the magnitude of the value function changes, and therefore the ranks of the value functions are all either rank 7 or 8. Changing the bounds of the control space does not greatly affect the ranks of this problem.

Next, we consider the effect of the diffusion magnitude on the optimal value function and its rank. The results of this experiment are shown in Figure 4. We observe that the diffusion is influential for determining the rank of the problem. One striking pattern seen in Figure 4 is that as the diffusion decreases, the blue region grows in size. The blue region indicates low cost, and it intuitively represents the area from which a system will not enter the boundary. In other words, the system is more controllable when there is less noise. When the noise is very large, for example in the upper left panel, there is a large chance that the Brownian motion can push the state into the boundary. This effect causes the terminal cost

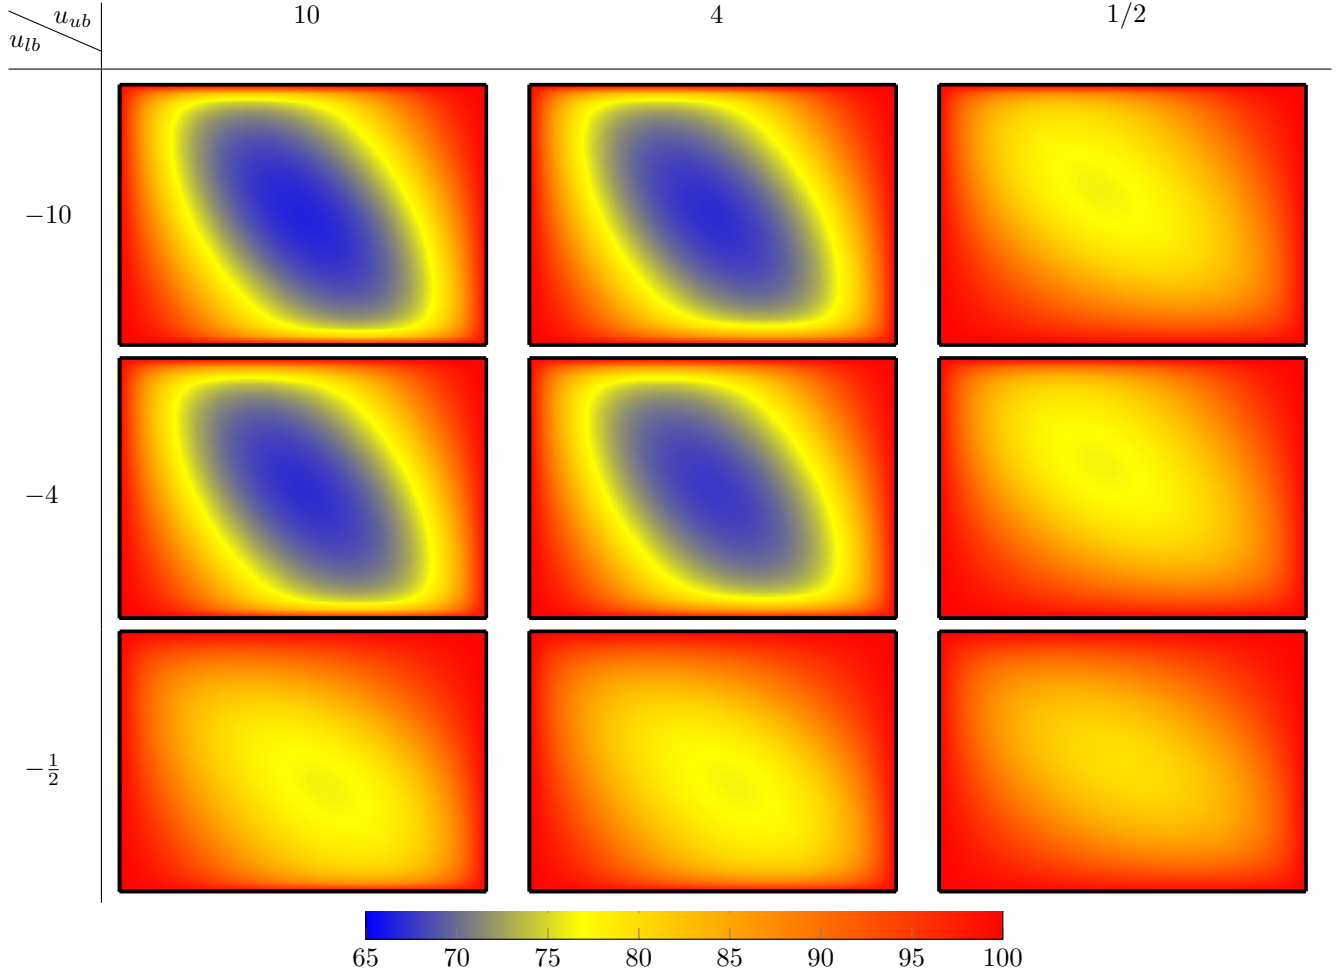


Fig. 3: Cost functions and ranks of the solution to the LQG problem for varying control bounds  $[u_{lb}, u_{ub}]$ . Diffusion magnitude is  $\sigma_1 = \sigma_2 = 1$ , and absorbing boundary conditions are used. The FT ranks found through the cross approximation algorithm with rounding tolerance  $\epsilon_{\text{round}} = 10^{-7}$  were either 7 or 8. The x-axis of each plot denotes  $x_1$  and the y-axis denotes  $x_2$ .

to propagate further into the interior of the domain. As  $\sigma_2$  decreases, there is less noise affecting the acceleration of the state, and the system remains controllable for a wide range of velocities. However, since the diffusion magnitude is large in the equation for velocity, the value function is only small when the position is close to the origin. In the area close to the origin there is less of a chance for the state to be randomly pushed into the absorbing region. Finally, when both diffusions are small, the system is controllable from a far greater range of states, as indicated by the lower right panel.

The ranks of the value functions follow the same pattern as controllability. The ranks are small when the diffusion terms are large, and high then the diffusion terms are small. When the features of the function are aligned with the coordinate axes, the ranks remain low. As the function becomes more complex due to small diffusion terms, the boundaries between guaranteed absorption and nonabsorption begin to have more complex shapes, resulting in increased ranks.

Next, we investigate the value functions associated with reflecting boundary conditions. The results of this experiment are shown in Figure 5. The results indicate that neither that

value functions nor their ranks are much affected by the bounds of the control space. Notably, the value functions in all examples are of rank 3, and they all appear to be close to a quadratic function. We note that, for a classical LQR problem the solution is quadratic, and therefore also has a value function with rank 3. Thus, in some sense, reflecting boundary conditions more accurately represent a the classical problem with linear dynamics, quadratic cost, and Gaussian process noise, but with no state or input constraints.

Next, we consider the effect of the magnitude of the diffusion terms on the optimal value function and its rank, this time in problem instances involving reflecting boundary conditions. The results of this experiment are shown in Figure 6, where it is evident that diffusion magnitude is again influential for determining the rank of this problem. The shapes and ranks of these value functions are similar to those in the case of absorbing boundary conditions. They follow the same pattern of increasing rank as the diffusion magnitude decreases.

2) *Convergence*: In this section, we focus on the convergence properties of the proposed algorithms in computational experiments. We discuss convergence in terms of (i) the norm

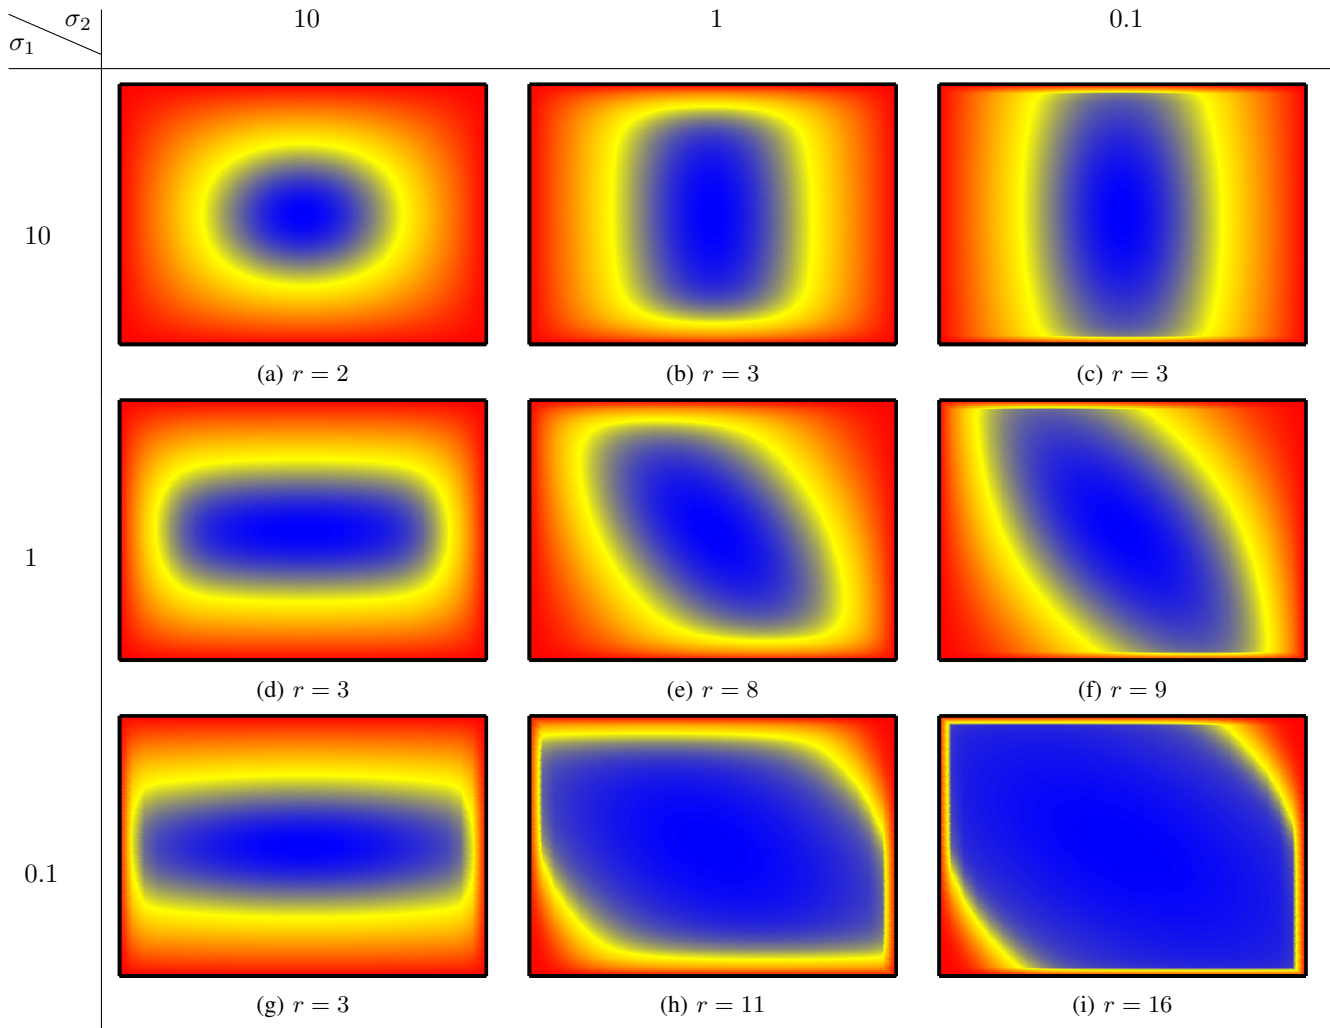


Fig. 4: Cost functions and ranks of the solution to the LQG problem for different  $\sigma_1, \sigma_2$ . We fix  $u_{lb} = -3$ ,  $u_{ub} = 3$ , and use absorbing boundary conditions. The FT ranks found through the cross approximation algorithm with rounding tolerance  $\epsilon_{\text{round}} = 10^{-7}$  are indicated by the caption for each frame. The x-axis of each plot denotes  $x_1$  and the y-axis denotes  $x_2$ . The color scale is different in each plot to highlight each function's shape.

of the value function, (ii) the difference between iterates, and (iii) the fraction of states visited during each iteration. We consider both the FT-based value iteration and the FT-based policy iteration algorithms.

Let us first consider the FT-based value iteration given in Algorithm 3. We use FT tolerances of  $\delta_{\text{cross}} = \epsilon_{\text{round}} = 10^{-7}$  as input to the algorithm. In Figure 7, we compare the convergence and the computational cost for solving the stochastic optimal control problem for varying discretization of the Markov chain approximation method, specifically discretization sizes of  $n = 25$ ,  $n = 50$ , and  $n = 100$  points along each dimension. Note that this corresponds to  $25^2 = 625$ ,  $50^2 = 2,500$ , and  $100^2 = 10,000$  total number of discrete states, respectively.

The results demonstrate that approximately the same value function norm is obtained regardless of discretization. However, convergence is much faster for coarse discretizations, hence the one-way multigrid algorithm may be useful. Furthermore, we observe that the low-rank nature of the problem emerges because the fraction of discretized states evaluated

during cross approximation for each iteration decreases with increasing grid resolution. For reference, the standard value iteration algorithm would have the fraction of state space evaluated be 1 for each iteration, as the standard value iteration would evaluate all discrete states. Thus, even in this two-dimensional example, the low-rank algorithm achieves between two to five times computational gains for each iteration, when compared to the standard value iteration algorithm.

Next, we repeat this experiment for FT-based policy iteration algorithm given by Algorithm 4. We use 10 sub-iterations to solve for the value function for every policy, and we use FT tolerances of  $\delta_{\text{cross}} = \epsilon_{\text{round}} = 10^{-7}$  as input to the algorithm. Figure 8 shows the results for the value function associated with each control update. The sub-iteration cost functions are not plotted. Note that, as expected, far fewer iterations are required for convergence. We observe similar solution quality when compared to the FT-based value iteration algorithm; however, we observe that convergence occurs using almost an order of magnitude fewer number of iterations.



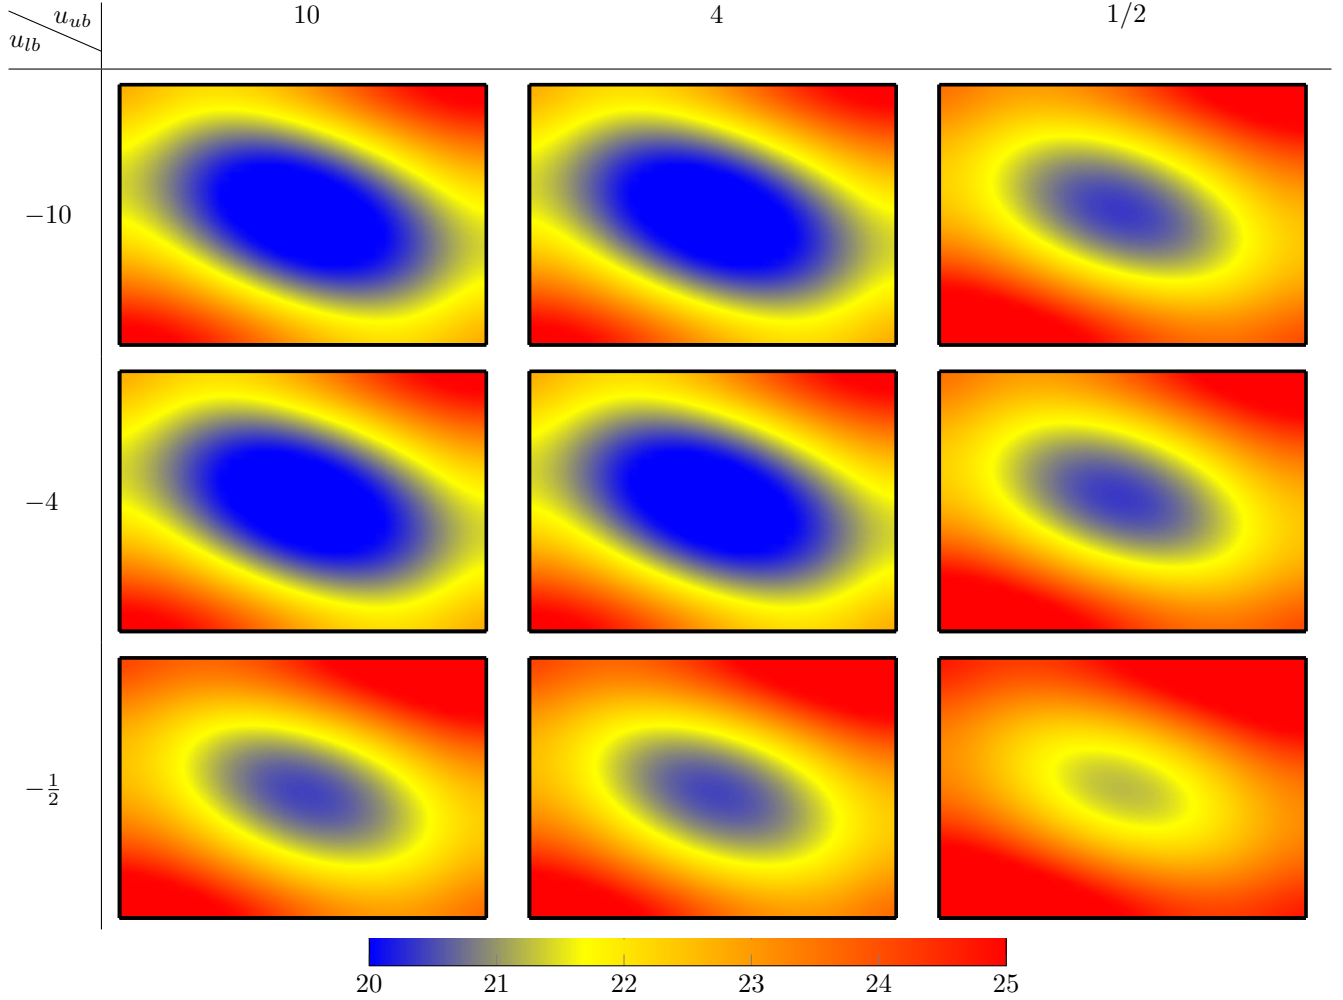


Fig. 5: Reflecting boundary conditions with cost functions and ranks for different  $[u_{lb}, u_{ub}]$ , and the diffusion is set to  $\sigma_1 = \sigma_2 = 1$ . The FT ranks found through the cross approximation algorithm with rounding tolerance  $\epsilon_{\text{round}} = 10^{-7}$  were 3 for all cases.

To summarize, from these experiments we observe that the following: (i) the rank depends on the intrinsic complexity of the value function, which seems to increase with decreasing process noise when the state constraints are active; (ii) in problems with reflective boundary conditions, modeling problems similar to those without state constraints, the rank does not seem to change with varying process noise; (iii) even in these two dimensional problems we obtain substantial computational gains: the proposed algorithms evaluate two to five times less number of states to reach a high quality solution, when compared to standard dynamic programming algorithms; (iv) we observe that the FT-based policy iteration algorithm converges almost an order of magnitude faster than the FT-based value iteration algorithm in these examples.

#### B. Car-like robots maneuvering in minimum time

Next, we consider two examples, each modeling car-like robots. The first example is a standard Dubins vehicle. The Dubins vehicle dynamics is nonlinear, nonholonomic, and has a three-dimensional state space. The second example extends

the Dubins vehicle dynamics to vary speed; and at high speeds the vehicle understeers. The resulting system is also nonlinear and nonholonomic, and the state space is four dimensional.

For all of the examples in this section, we use the FT-based policy iteration algorithm given by Algorithm 4 with  $n_{fp} = 10$ , cross approximation and rounding tolerances set to  $10^{-5}$ .

1) *Dubins vehicle*: Consider the following dynamics:

$$\begin{aligned} dx &= \cos(\theta)dt + dw_1(t) \\ dy &= \sin(\theta)dt + dw_2(t) \\ d\theta &= u(t)dt + 10^{-2}dw_3(t), \end{aligned}$$

which is the standard Dubins vehicle dynamics with process noise added to each state. Consider bounded state space:  $x \in (-4, 4)$ ,  $y \in (-4, 4)$ , and  $\theta \in [-\pi, \pi]$ . The control space consists of three points  $\mathcal{U} = \{-1, 0, 1\}$ . Boundary conditions are absorbing for the position dimensions  $(x, y)$  and periodic for the angle  $\theta$ . An absorbing region is specified at the origin with a width of 0.5, and the terminal costs are

$$\psi(x, y, \theta) = \begin{cases} 0 & \text{for } (x, y) \in [-0.25, 0.25]^2 \\ 10 & \text{for } |x| \geq 2 \text{ or } |y| \geq 2 \end{cases}$$



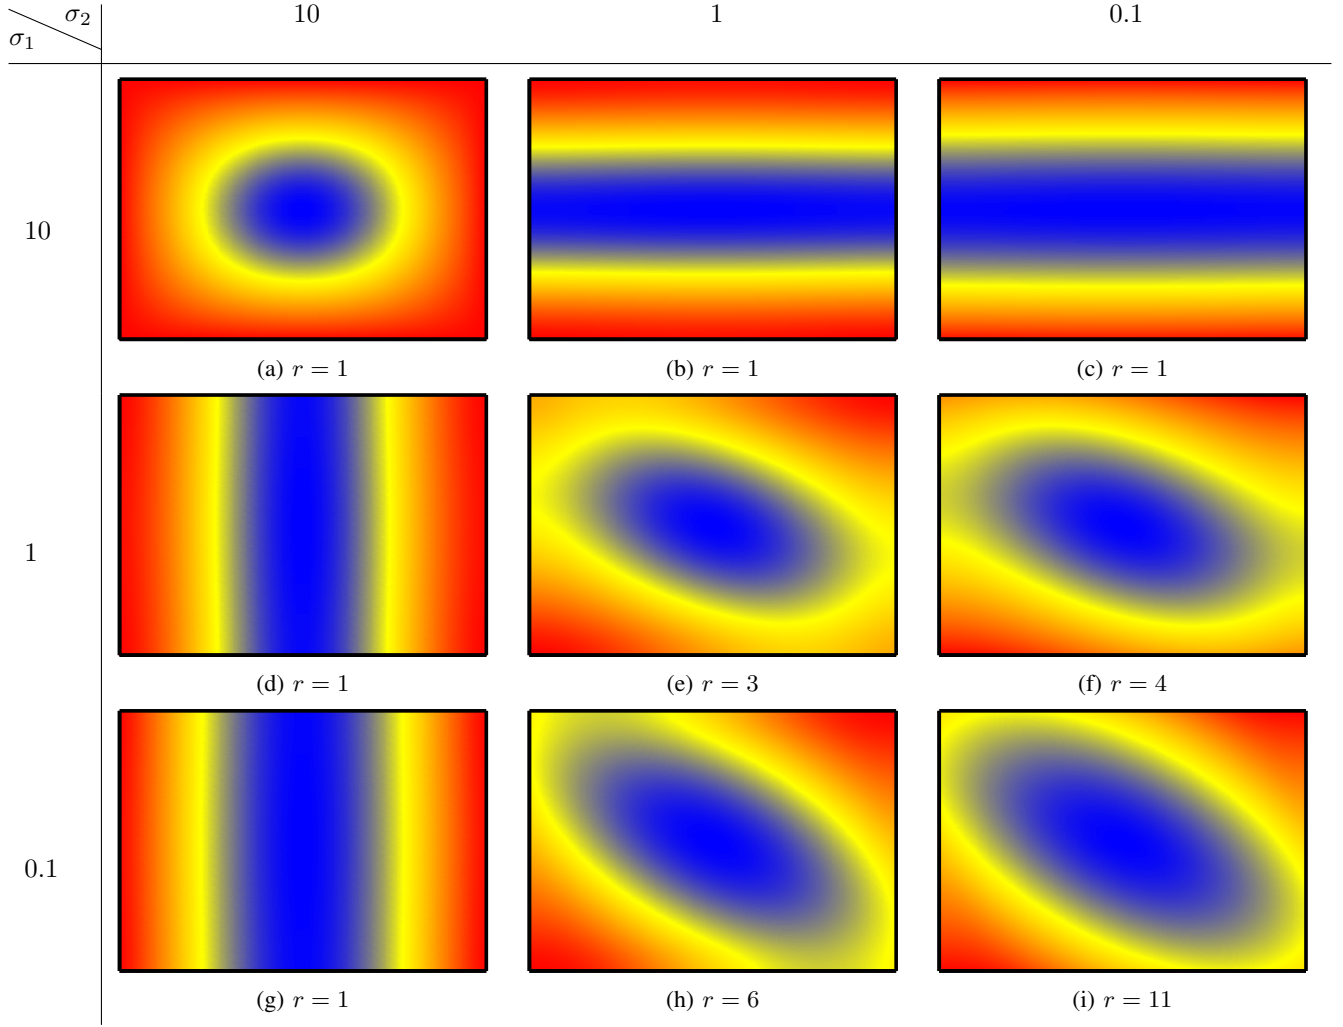


Fig. 6: Cost functions and ranks of the solution to the LQG problem for different  $\sigma_1, \sigma_2$ . We fix  $u_{lb} = -3$ ,  $u_{ub} = 3$ , and use reflecting boundary conditions. The FT ranks found through the cross approximation algorithm with rounding tolerance  $\epsilon_{\text{round}} = 10^{-7}$  are indicated by the caption for each frame. The x-axis of each plot denotes  $x_1$  and the y-axis denotes  $x_2$ . The color scale is different in each plot to highlight each function's shape.

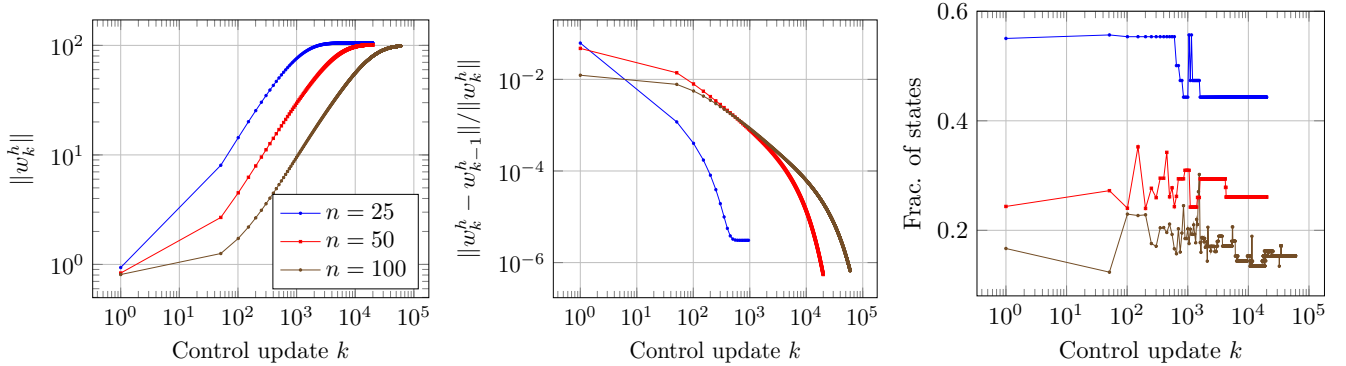


Fig. 7: FT-based value iteration diagnostic plots for the linear quadratic problem with reflecting boundaries,  $u(t) \in [-1, 1]$ , and  $\sigma_1 = \sigma_2 = 1$ . The left panel shows that for all discretization levels the value function norm converges to approximately the same value. The middle panel shows the relative difference between value functions of sequential iterations. The right panel shows that the fraction of states evaluated within cross approximation decreases with increasing discretization resolution.

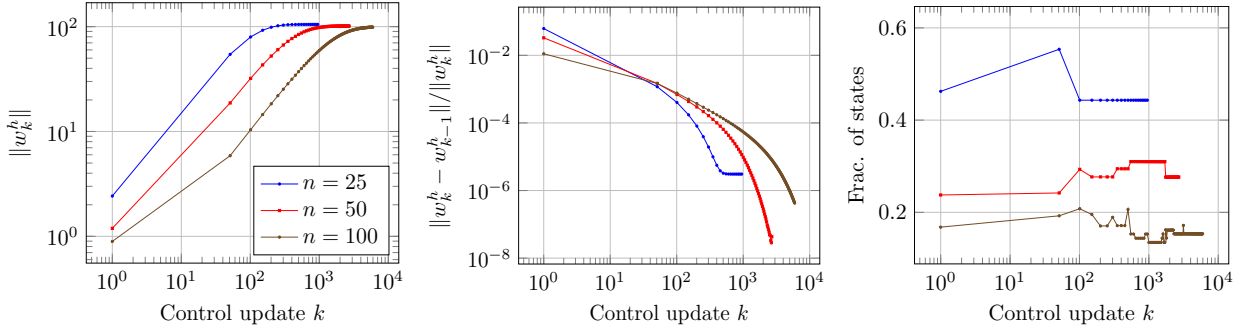


Fig. 8: FT-based policy iteration diagnostic plots for the linear quadratic problem with reflecting boundaries,  $u(t) \in [-1, 1]$ , and  $\sigma_1 = \sigma_2 = 1$ . The left panel shows that for all discretization levels the value function norm converges to same value. The middle panel shows the relative difference between value functions of sequential iterations. The right panel shows that the fraction of states evaluated decreases with increasing discretization.

To compute a time-optimal control, the stage costs is set to

$$g(x, u) = 1.$$

The solution is obtained using one-way multigrid as discussed in Section III-C5. First, one hundred steps of FT-based policy iteration are used with the each dimension discretized into  $n = 25$  points. Then, the result is interpolated onto a grid that is obtained by discretizing each dimension into and the problem is solved with fifty steps of FT-based policy iteration. This multigrid procedure is repeated for  $n = 100$  and  $n = 200$ .

A simulation of the resulting feedback controller for several initial conditions is shown in Figure 9.

Convergence plots are shown in Figure 10. It is worth noting at this point that these plots demonstrate one of the advantages of the FT framework: since the value function is represented as function rather than an array, we can compare the norms of functions represented with different discretizations. In fact, the upper left panel of Figure 10 demonstrates that the norm *continuously* decreases when the discretization is refined.

Furthermore, these results suggest that the solution to this problem, with an accuracy due to rounding of  $\epsilon_{\text{round}} = 10^{-5}$ , indeed has FT rank 12. This suggestion is supported by the fact that once the grid is refined enough, the fraction of states evaluated decreases, but the maximum rank levels off at 12. Note that, the total number of states changes throughout the iterations depending on the discretization level, i.e., for  $n = 25, 50, 100, 200$ , we have that the total number of states is  $25^3 = 15,625$ ,  $50^3 = 125,000$ ,  $100^3 = 1,000,000$  and  $200^3 = 8,000,000$ . That is, at  $n = 200$ , the total number of states reaches 8 million.

The lower right panel shows the fraction of states evaluated at various iterations. We observe that when  $n = 100$ , we evaluate only 10% of the states in each iteration. This number improves when  $n = 200$ . Hence, the FT-based algorithm provides an order of magnitude computational savings in terms of number of states evaluated, when compared to standard dynamic programming algorithms, even in this three-dimensional problem.

2) *Understeered car*: In this section, we consider a model that extends the Dubins vehicle. In this model, the speed of the vehicle is another state that can be controlled. If the speed

is larger the vehicle starts to understeer, modeling skidding behavior. The states  $(x, y, \theta, v)$  are now the  $x$ -position, the  $y$ -position, orientation, and velocity, respectively. The control for steering angle is  $u_1(t) \in [-15\frac{\pi}{180}, 15\frac{\pi}{180}]$  and for acceleration is  $u_2(t) \in [-1, 1]$ . The dynamical system is described by

$$\begin{aligned} dx &= v \cos(\theta)dt + dw_1(t) \\ dy &= v \sin(\theta)dt + dw_2(t) \\ d\theta &= \frac{1}{1 + (v/v_c)L} \frac{v}{L} \tan(u_1(t))dt + 10^{-2}dw_3(t) \\ dv &= \alpha u_2(t)dt + 10^{-2}dw_4(t) \end{aligned}$$

where  $v_c = 8$  m/s is the characteristic speed,  $L = 0.2$  m is the length of the car, and  $\alpha = 2$  is a speed control constant. The boundary conditions are absorbing for the the positions  $x$  and  $y$ , periodic for  $\theta$ , and reflecting for  $v$ . The space for the positions and orientations are identical to that of the Dubins vehicle. The velocity is restricted to forward with  $v \in [3, 5]$ . The stage cost is altered to push the car to the center

$$g(x, y, \theta, v) = 1 + x^2 + y^2.$$

The terminal cost is the same as for the Dubins vehicle with an absorbing region at the origin of the  $x, y$  plane.

We remark that the dynamics of this example are *not* affine in control input. Therefore, this example does not fit into many standard frameworks that solve a corresponding linear HJB equation [29], [62].

The trajectories for various starting locations are shown in Figure 11. Figure 12 shows the convergence plots. Due to computational considerations, we fixed a maximum FT adaptation rank to 20. Therefore, instead of plotting the maximum rank, we plot the *average* FT rank in the lower left panel. The average rank varies more widely for coarse discretizations, than for fine discretizations. But, when  $n = 100$ , the average rank becomes smaller and more consistent.

Let us note that, when  $n = 100$ , the number of discrete states is  $100^4 = 100,000,000$ , i.e., 100 million. At this discretization level, the proposed FT-based algorithm evaluates less than 1% of the states, leading to more than two orders of magnitude computational savings when compared to the standard dynamic programming algorithms.

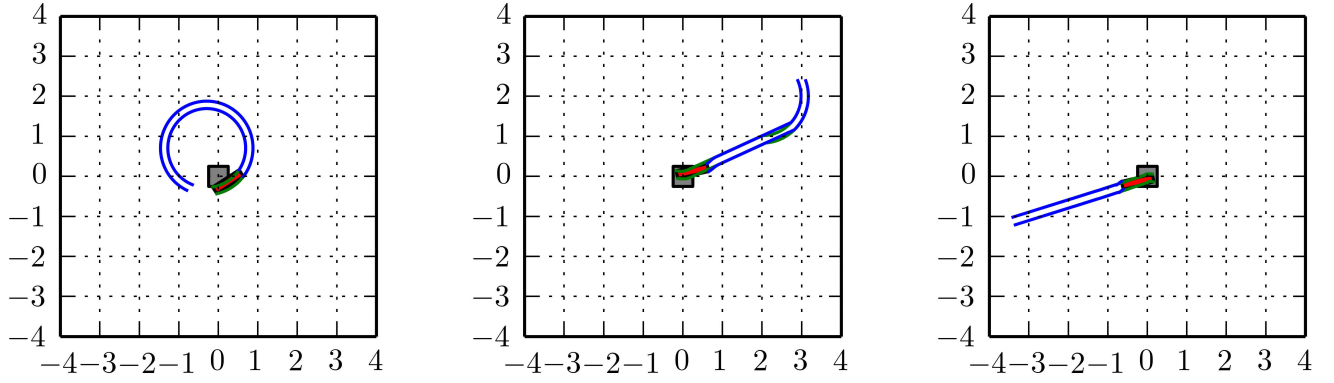


Fig. 9: Trajectories of the Dubins car for three initial conditions. Panels show the car when it arrives in the absorbing region.

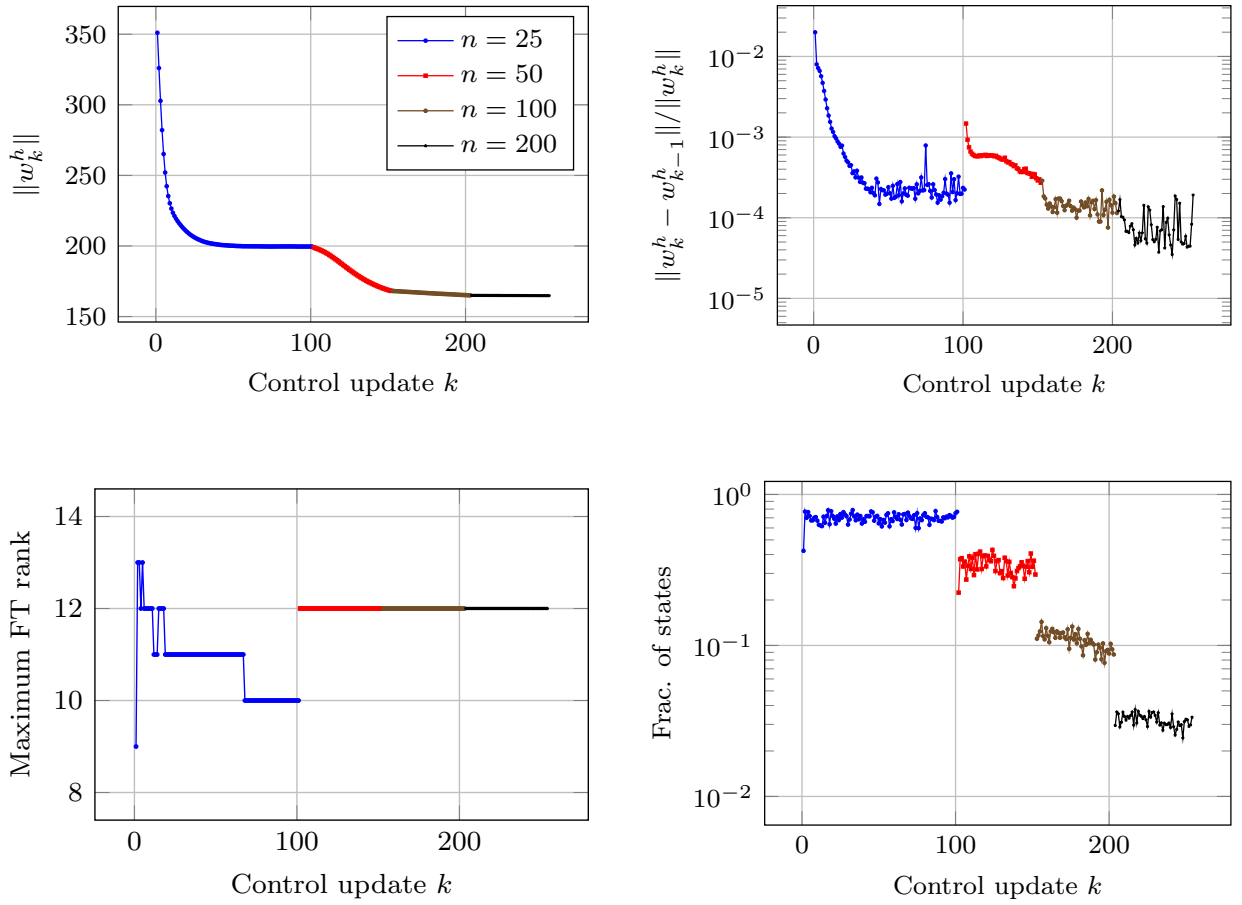


Fig. 10: One-way multigrid for solving the Dubins car control problem.

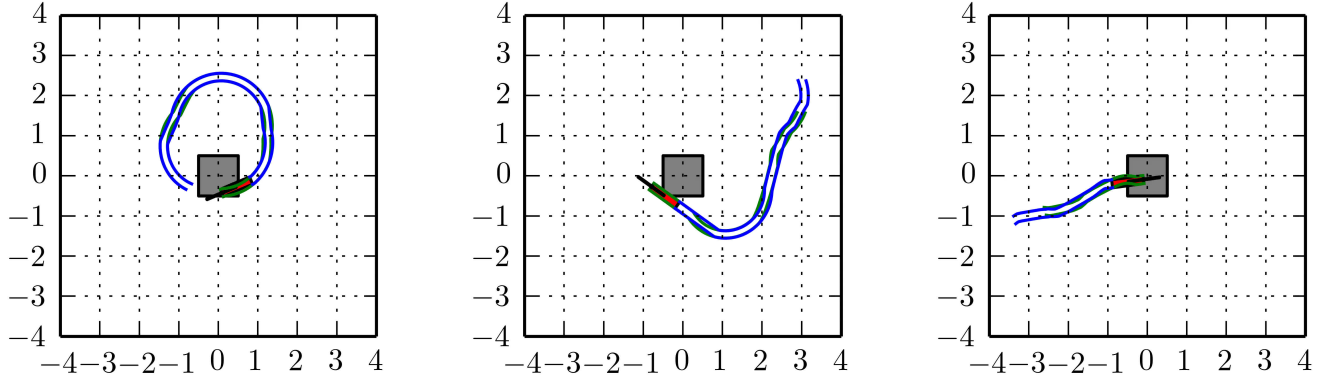


Fig. 11: Trajectories of the understeered car for three initial conditions. Panels show the car when it arrives in the absorbing region.

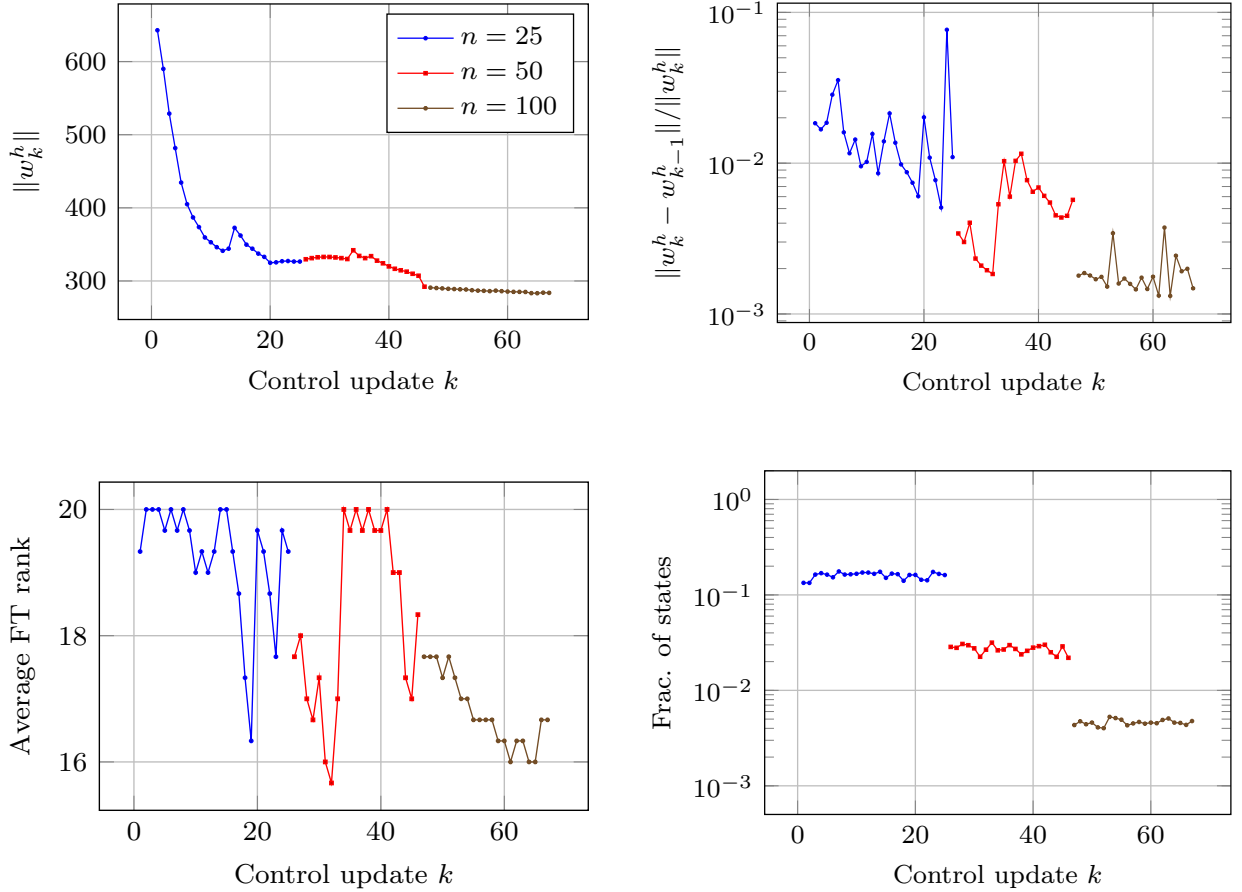


Fig. 12: One-way multigrid for solving the understeered car control problem. Maximum rank FT rank is restricted to 20.

It is worth noting that, even storing the optimal control as a standard lookup table in memory would require a large storage space, if the controller was computed using standard dynamic programming algorithms. At  $n = 100$ , the storage required is around  $0.8 \times 10^9$  bytes, or 800 MB, assuming we are storing floating point values. The proposed algorithm naturally stores the controller in a compressed format that leads to two orders of magnitude savings in storage as expected from the fraction of states evaluated. In fact, storing the final value function in the FT format required almost exactly 1MB of storage for this experiment. These savings can be tremendously beneficial in constrained computing environments. For example, these controllers can potentially be used on embedded systems that have serious memory constraints.

### C. Glider perching on a string

We now consider a problem involving a glider with a seven-dimensional state space. The mission is to control the glider to perch on a horizontal string. This problem has been widely studied in the literature [14], [45], [55]. To the best of our knowledge, no optimal controller is known. We compute a controller using the FT-based dynamic programming algorithms

The glider is described by flat-plate model in the two-dimensional plane involving seven state variables, namely  $(x, y, \theta, \phi, v_x, v_y, \dot{\theta})$ , specifying its  $x$ -position,  $y$ -position, angle of attack, elevator angle, horizontal speed, vertical speed, and the rate of change of the angle of attack, respectively. The input control is the rate of change of the elevator angle  $u = \dot{\phi}$ .

A successful perch is defined by a horizontal velocity between 0 and 2 m/s, a vertical velocity between -1 and -3 m/s, and the  $x$  and  $y$  positions of the glider within a 5cm radius of the perch. Under these conditions, Cory and Tedrake [14] has shown that the experimental aircraft can attach to the string. For more information on this experimental platform, the reader is referred to either Roberts et al. [55] or Moore and Tedrake [45]. The dynamics of the glider are given by

$$\begin{aligned} \mathbf{x}_w &= [x - l_w c_\theta, y - l_w s_\theta], \\ \dot{\mathbf{x}}_w &= [\dot{x} + l_w \dot{\theta} s_\theta, \dot{y} - l_w \dot{\theta} c_\theta] \\ \mathbf{x}_e &= [x - l c_\theta - l_e, c_{\theta+\phi}, y - l s_\theta - l_e s_{\theta+\phi}] \\ \dot{\mathbf{x}}_e &= [\dot{x} + l \dot{\theta} s_\theta + l_e (\dot{\theta} + u) s_{\theta+\phi}, \dot{y} - l \dot{\theta} c_\theta - l_e (\dot{\theta} + u) c_{\theta+\phi}] \\ \alpha_w &= \theta - \tan^{-1}(\dot{y}_w, \dot{x}_w), \quad \alpha_e = \theta + \phi - \tan^{-1}(\dot{y}_e, \dot{x}_e) \\ f_w &= \rho S_w |\dot{\mathbf{x}}_w|^2 \sin(\alpha_w), \quad f_e = \rho S_e |\dot{\mathbf{x}}_e|^2 \sin(\alpha_e) \end{aligned}$$

$$\begin{aligned} dx &= v_x dt + 10^{-9} dw_1(t) \\ dy &= v_y dt + 10^{-9} dw_2(t) \\ d\theta &= \dot{\theta} dt + 10^{-9} dw_3(t) \\ d\phi &= u dt + 10^{-9} dw_4(t) \\ dv_x &= \frac{1}{m} (-f_w s_\theta - f_e s_{\theta+\phi}) dt + 10^{-9} dw_5(t) \\ dv_y &= \frac{1}{m} (f_w c_\theta + f_e c_{\theta+\phi} - mg) dt + 10^{-9} dw_6(t) \\ d\dot{\theta} &= \frac{1}{I} (-f_w l_w - f_e (l c_\phi + l_e)) dt + 10^{-9} dw_7(t) \end{aligned}$$

where  $\rho$  is the density of air,  $m$  is the mass of the glider,  $I$  is the moment of inertia of the glider,  $S_w$  and  $S_e$  are the

surface areas of the wing and tail control surfaces,  $l$  is the length from the center of gravity to the elevator,  $l_w$  is the half chord of the wing,  $l_e$  is the half chord of the elevator,  $c_\gamma$  denotes  $\cos(\gamma)$ , and  $s_\gamma$  denotes  $\sin(\gamma)$ . The values of these parameters are chosen to be the same as those proposed by Roberts et al. [55].

Absorbing boundary conditions are used, and an absorbing region is defined to encourage a successful perch. Let this region be defined for  $x \in [-0.05, 0.05]$ ,  $y \in [-0.05, 0.05]$ ,  $v_x \in [-0.25, 0.25]$ , and  $v_y \in [-2.25, 2.25]$ .

The stage cost is specified as

$$g(x, y, \theta, \phi, v_x, v_y, \dot{\theta}) = 20x^2 + 50y^2 + \phi^2 + 11v_x^2 + v_y^2 + \dot{\theta}^2.$$

The terminal cost for both of these regions is

$$\psi(x, y, \theta, \phi, v_x, v_y, \dot{\theta}) = \begin{cases} 0 & \text{if perched} \\ 600x^2 + 400y^2 + \frac{1}{9}\theta^2 + \frac{1}{9}\phi^2 + v_x^2 + (v_y + 1.5)^2 + \frac{1}{9}(\dot{\theta} + 0.5^2) & \text{otherwise.} \end{cases}$$

A controller is computed using the FT-based one-way multigrid algorithm. Several trajectories of the controlled system under this controller are shown in Figure 13. These trajectories are generated by starting the system from different initial states. Notice that they all follow the same pattern: first dive, then climb, and finally drop into the perch. This behavior is similar to that demonstrated in experiments by Cory and Tedrake [14], Roberts et al. [55], and Moore and Tedrake [45].

Figure 14 shows the convergence diagnostic plots for the one-way multigrid algorithm used for solving this problem. We used discretization levels of  $n = 20, 40, 80, 160$  points along each dimension. Thus, the number of discretized states at the finest discretization is  $160^7 \approx 2.6 \times 10^{15}$ , i.e., 2.6 quadrillion or 2,600,000 billion. Furthermore, we limited the rank to a maximum of 10, and the panel in lower left panel shows that the average ranks are below this maximum threshold. In terms of the number of states evaluated, the final ranks of the value function translate into savings of approximately four orders of magnitude *per iteration* when  $n = 20$ , and a corresponding savings of ten orders of magnitude *per iteration* when  $n = 160$ . However, in this example it seems that the rank truncation is indeed affecting the convergence of the problem. A noisy and volatile decay of the value function norm is seen in the upper left panel. Furthermore, the difference between iterates, in the upper right panel, indicates a relative error of approximation  $10^{-2}$  which is above our FT rounding threshold of  $10^{-5}$ . Nonetheless, the resulting controller seems to be successful at achieving the desired behavior. The storage size for the controller is stored in 940 kB, or approximately 1MB. This means that for an  $n = 80$  controller, the resulting storage cost would require  $\mathcal{O}(10^8)$  MB or  $\mathcal{O}(100)$  TB.

### D. Quadcopter flying through a window

Finally, we consider the problem of maneuvering a quadcopter through a small target region, such as a window. The resulting dynamical system modeling the quadcopter has six states and three controls. The states  $(x, y, z, v_x, v_y, v_z)$  are

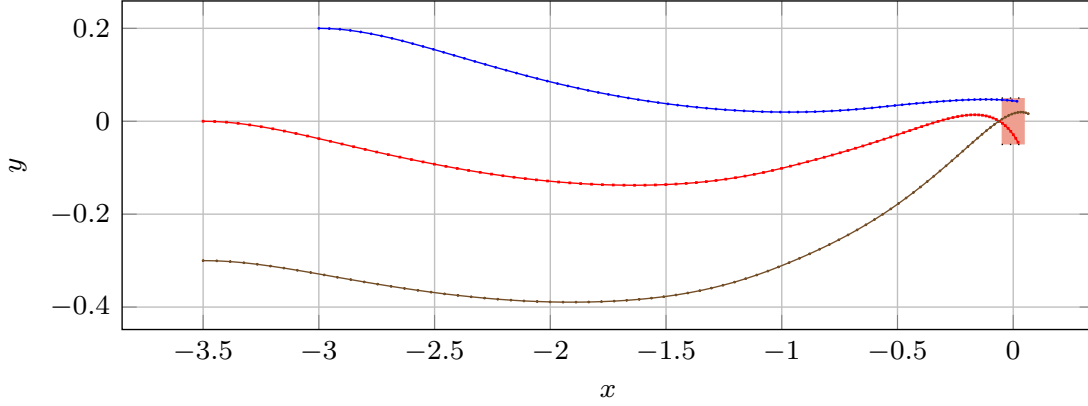


Fig. 13: Trajectories of the perching glider for three initial conditions given by:  $(-3, \frac{1}{5}, 0, 0, \frac{9}{2}, 0, 0)$  (blue),  $(-\frac{7}{2}, 0, 0, 0, 5, 0, 0)$  (red), and  $(-\frac{7}{2}, -\frac{3}{10}, 0, 0, 5.8, 0, 0)$  (brown). Target region is shown by shaded region that is centered at  $(0, 0)$

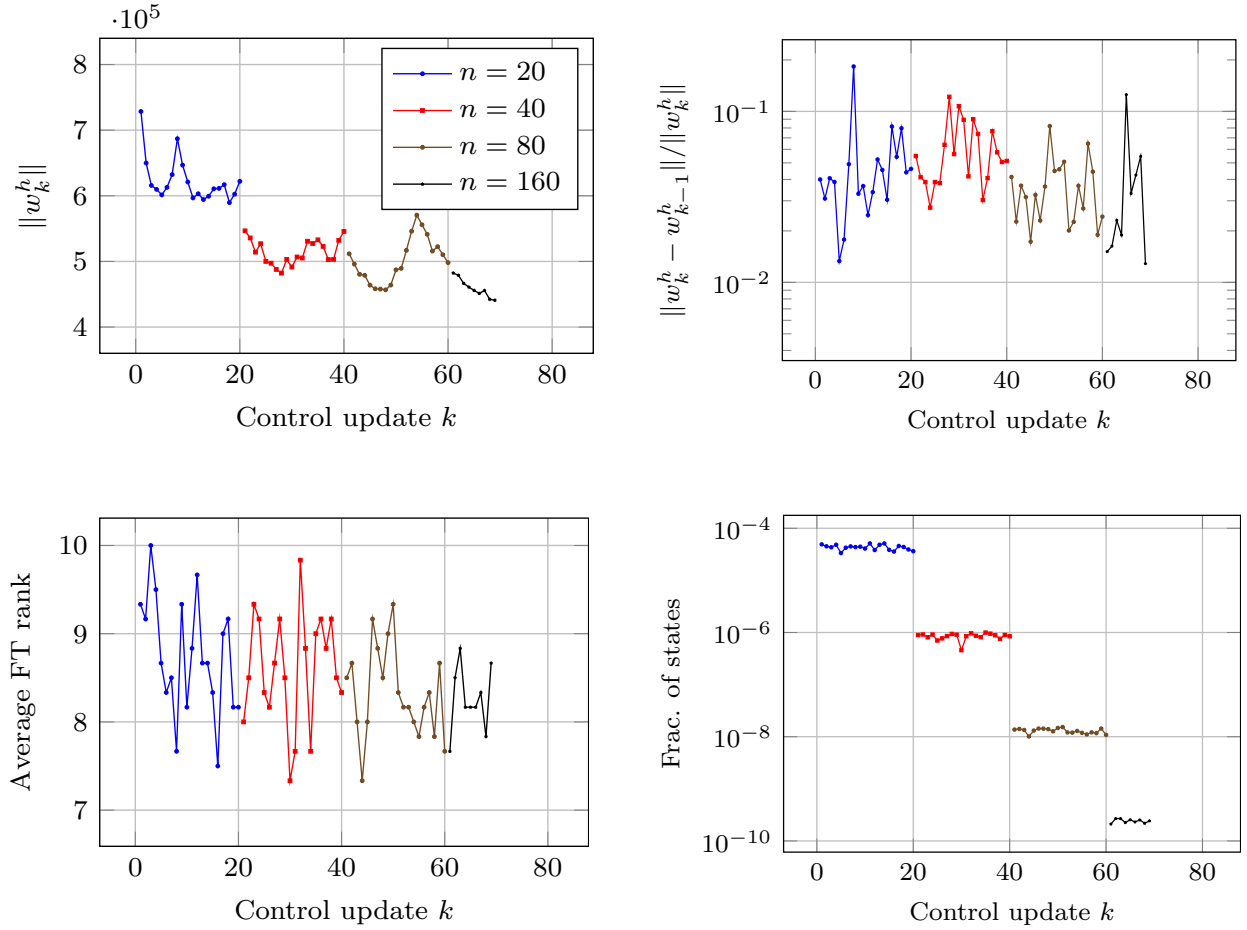


Fig. 14: One-way multigrid for solving perching glider. Maximum rank FT rank is restricted to 10.

the  $x$ -position in meters  $x \in [-3.5, 3.5]$ ,  $y$ -position in meters  $y \in [-3.5, 3.5]$ ,  $z$ -position in meters  $z \in [-2, 2]$ ,  $x$ -velocity in meters per second  $v_x \in [-5, 5]$ ,  $y$ -velocity in meters per second  $v_y \in [-5, 5]$ , and  $z$ -velocity in meters per second  $v_z \in [-5, 5]$ . The controls are the thrust (offset by gravity)  $u_1 \in [-1.5, 1.5]$ , the roll angle  $u_2 \in [-0.4, 0.4]$ , and the pitch angle  $u_3 \in [-0.4, 0.4]$ . Then, the dynamical system is described by the following stochastic differential equation:

$$\begin{aligned} dx &= v_x dt + 10^{-1} dw_1(t) \\ dy &= v_y dt + 10^{-1} dw_2(t) \\ dz &= v_z dt + 10^{-1} dw_3(t) \\ dv_x &= \frac{u_1 - mg}{m} \cos(u_2) \sin(u_3) dt + 1.2 dw_4(t) \\ dv_y &= -\frac{u_1 - mg}{m} \sin(u_2) dt + 1.2 dw_5(t) \\ dv_z &= \left( \cos(u_2) \cos(u_3) \frac{u_1 - mg}{m} + g \right) dt + 1.2 dw_6(t), \end{aligned}$$

and reflecting boundary conditions are used for every state. A similar model was used by Carrillo et al. [11].

A target region is specified as a cube centered at the origin, and a successful maneuver is one which enters the cube with a forward velocity of one meter per second with less than 0.15m/s speed in the  $y$  and  $z$  directions.

The terminal cost is assigned to be zero for this region, i.e.,

$$\psi(x, y, z, v_x, v_y, v_z) = 0,$$

for  $(x, y, z, v_x, v_y, v_z) \in [-0.2, 0.2]^3 \times [0.5, 1.5] \times [-0.2, 0.2]^2$ . The stage cost is set to

$$g(x, y, z, v_x, v_y, v_z, u_1, u_2, u_3) = 60 + 8x^2 + 6y^2 + 8z^2 + 2u_1^2 + u_2^2 + 6u_3^2.$$

The maximum FT-rank is restricted to 10. The tolerances were set as  $\delta_{\text{cross}} = \epsilon_{\text{round}} = 10^{-5}$ . While, theoretically, underestimating the ranks can potentially cause significant approximation errors, in this case we are still able to achieve a well performing controller. Investigating the effect of rank underestimation is an important area of future work.

Trajectories of the optimal controller for various initial conditions are shown in Figures 15 and 16. In Figure 15, the position and velocities each approach their respective absorbing conditions for each simulation. In the first simulation, the quadcopter quickly accelerates and decelerates into the goal region. The final positions do not lie exactly within the absorption region. This absorption region is, in a way, unstable since it requires the quadcopter enter with a forward velocity. The forward velocity requirement virtually guarantees that the quadcopter must eventually exit the absorption region. The second simulation starts with positive  $y$  and  $z$  velocities. The third simulation starts with a negative  $y$  velocity and the quadcopter far away from the origin.

Note that the controls are all fairly smooth except when the quadcopter state is near or in the absorption region. At this point, the roll angle (and pitch angle in the first and third simulations) oscillates rapidly around zero. We conjecture this behavior is due to “flatness” of the objective function. Since



Fig. 18: The quadcopter utilized in the experiments.

the quadcopter is so close to the absorption region the control inputs can change rapidly to ensure it stays there.

Figure 17 shows the convergence diagnostics. We used a one-way multigrid strategy with  $n = 20$ ,  $n = 60$ , and  $n = 120$  discretization nodes. The difference between iterates, shown in the upper right panel, is between 1 and 0.1. In fact, this means that the relative difference is approximately  $10^{-4}$ , which matches well with the algorithm tolerances  $\epsilon = (\delta_{\text{cross}}, \epsilon_{\text{round}})$ .

The lower right panel indicates computational savings between three and seven orders of magnitude for each iteration, in terms of the number of states evaluated. In terms of storage space, if we had used the standard value iteration algorithms, storing the value function as a complete lookup table for  $n = 120$  would require storing  $10^6 = 2 \times 10^{12}$  floating point numbers, which is approximately 24 TB of data. The value function computed using the proposed algorithms required only 778 KB of storage space.

### E. Experiments with a quadcopter in a motion capture room

In this section, we report the results of experiments involving a quadcopter. We utilize a motion capture system for state estimation. We compute the controller offline, as described in the previous section. But, we run the resulting controller on a computer on board the quadcopter in real time. The experimental setup, the real-time control execution, and the experimental results are detailed below.

1) *Quadcopter hardware:* We use a custom-build quadcopter vehicle shown in Figure 18. The vehicle is equipped with an embedded computer, namely an Nvidia Tegra K1. An Arduino Nano computer controls the propellor motors, and it is connected to the embedded computer via the USB bus. The drone also includes a camera and an inertial measurement unit, which were not utilized in this experiment.

2) *Motion capture system:* We utilize an OptiTrack Motion Capture system<sup>6</sup> with six Flex 13 cameras. We place the cameras such that the pose (position and orientation) estimates are reliably obtained within a roughly 2-meter wide, 5-meter long and 2-meter high volume. The system provides pose estimates in 360Hz. We place passive infrared markers on the quadcopter to track its pose when it is in this volume. The

<sup>6</sup><http://www.optitrack.com/>



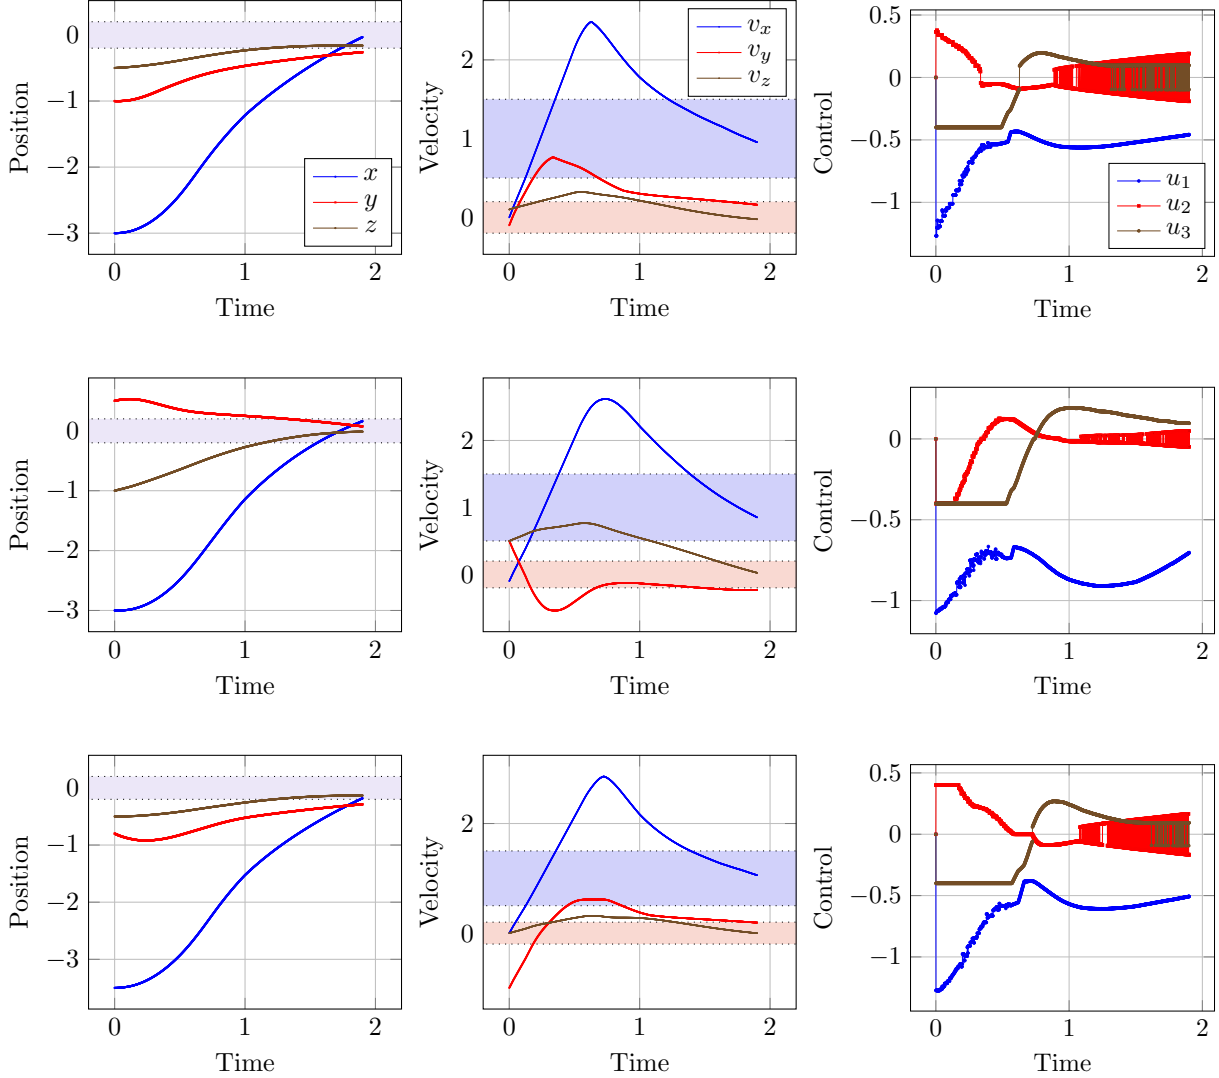


Fig. 15: Three quadcopter simulations using optimal low-rank feedback controller. Shaded region indicates the target positions and velocities.

pose information is captured by the motion capture computer and sent to the drone via a wifi link.

3) *Real-time execution of the the controller:* Once the vehicle gets the pose information via the wifi link, the control inputs, namely the four motor speeds, are computed using an Nvidia Tegra K1 computer that is on board the vehicle.

We utilize a cascade controller architecture. First, high-level controller determines the desired pitch angle, roll angle, and thrust from the vehicle pose. Then, a low-level controller commands the motor speeds to follow the desired pitch angle, roll angle, and thrust. The high-level controller is designed using the proposed algorithm as in Section VII-D. The low-level controller is a PD controller, described in [54].

The high-level controller execution follows exactly the same steps as for the simulated system and is described in the introduction to Section VII. It consists of two steps: (i) computing the value corresponding to the neighboring states of

the current state of the quadcopter, obtained from the motion capture system; (ii) obtaining the minimizer of Equation (4) through a multistart Newton-based BFGS optimization scheme within the  $C^3$  library. Note that the value function is computed offline and stored in the FT format, identically to the simulated results, and thus its evaluation at a particular state requires the multiplication of six matrices.

4) *Experimental results:* Figure 19 shows an image of the quadcopter at various times through its flight. Note that between the last two images, the quadcopter passes through the window. Figure 20 shows motion capture data for the quadcopter entering the goal region; the velocities are shown on the left panel, and the path is shown on the right. As seen from the figures, the vehicle is able to fly the through the window with the intended velocities.



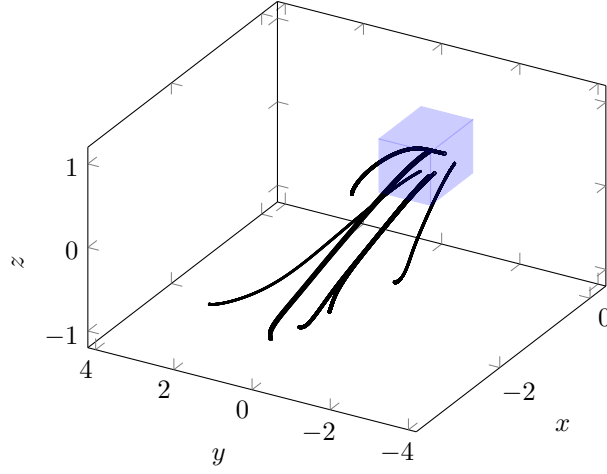


Fig. 16: Three-dimensional trajectories showing the position of the quadcopter for various initial conditions.

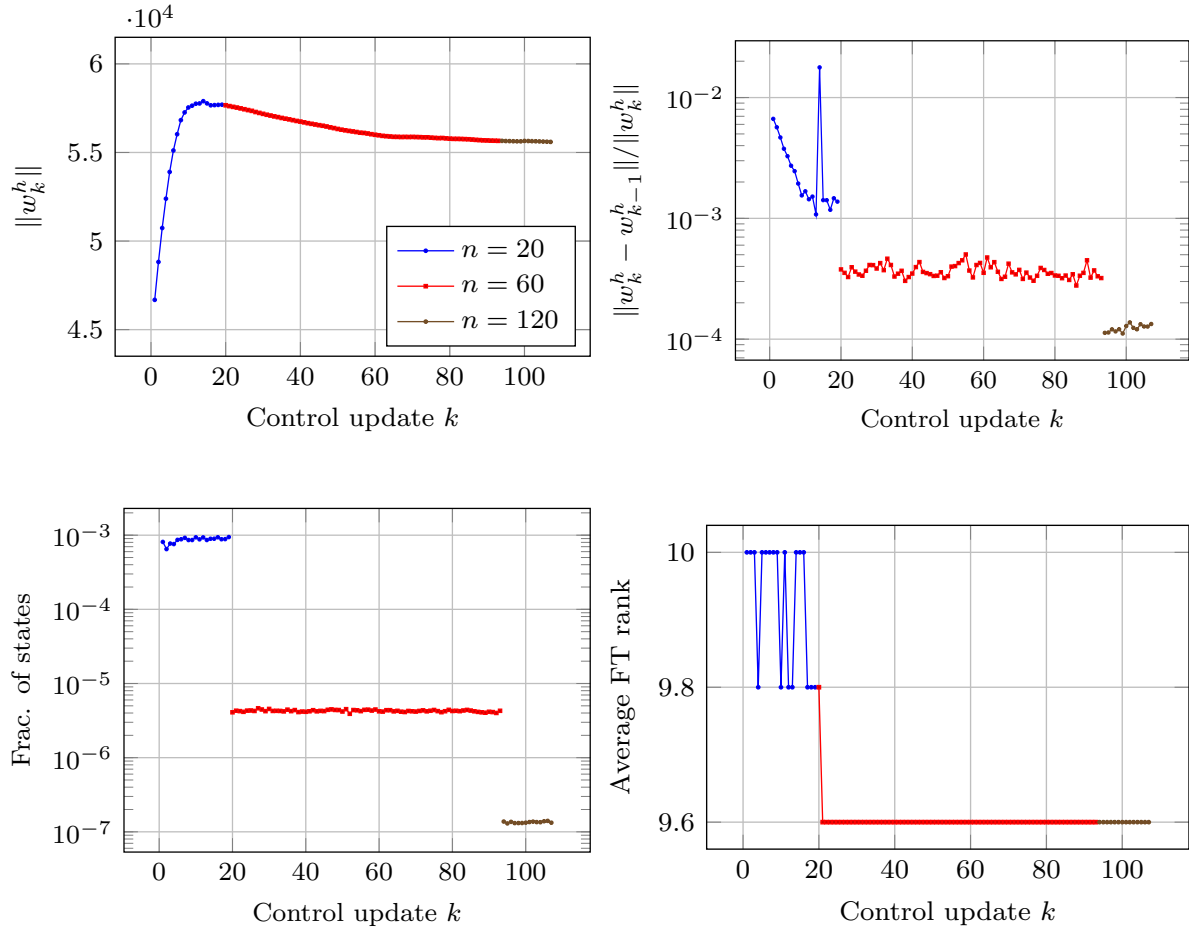


Fig. 17: One-way multigrid for solving the quadcopter control problem. Maximum rank FT rank is restricted to 10.

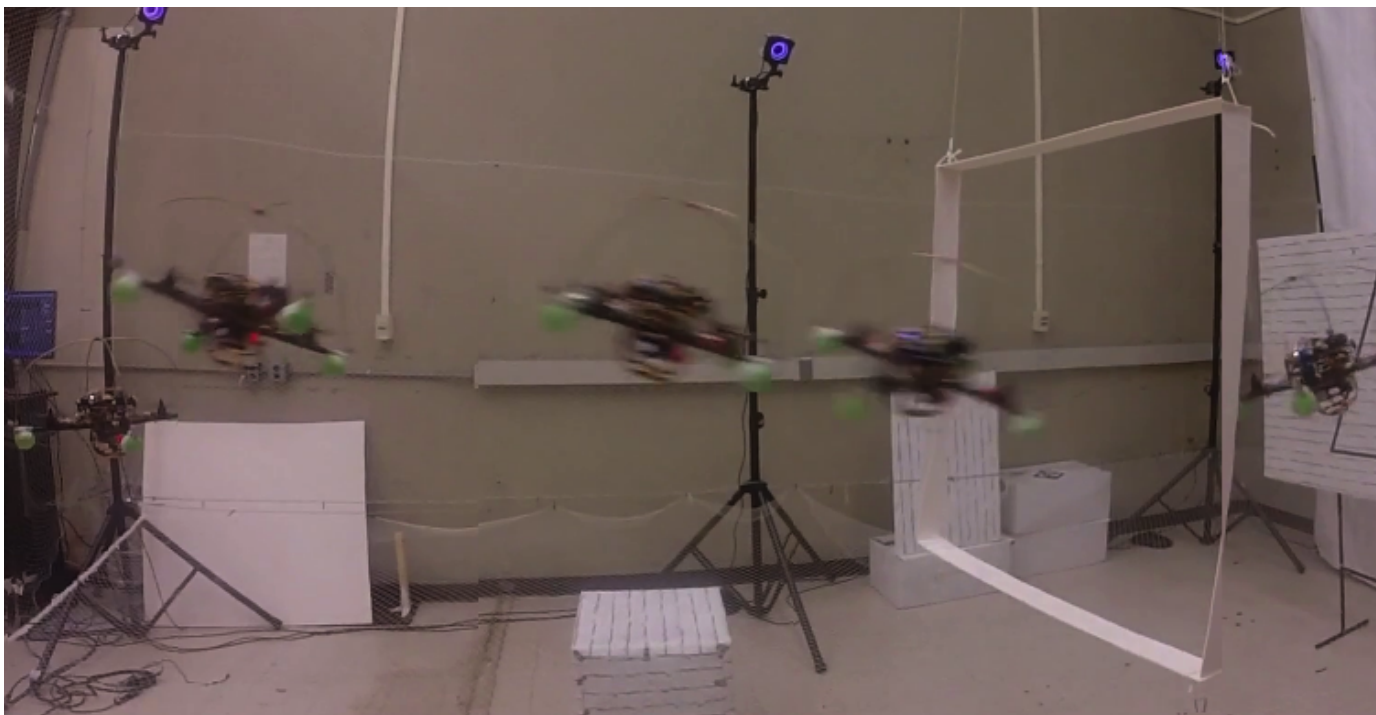


Fig. 19: Time lapse image showing the quadcopter passing through the window.

### VIII. CONCLUSION

In this paper, we proposed novel algorithms for dynamic programming, based on the compressed continuous computation framework using the function train (FT) decomposition algorithm. Specifically, we considered continuous-time continuous-space stochastic optimal control problems. We proposed FT-based value iteration, FT-based policy iteration, and an FT-based one-way multigrid algorithms. All algorithms represent the value function in the FT format, and they apply multilinear algebra operations in the (compressed) FT format. We analyzed the algorithms in terms of convergence and computational complexity. We have shown conditions under which the algorithms guarantee convergence to an optimal control. We have also shown that the algorithms scale polynomially with the dimensionality of the state space of the underlying stochastic optimal control problem and polynomially with the rank of the optimal value function. We have also demonstrated the proposed control synthesis algorithms on various problem instances. In particular, we have shown that the proposed algorithm can solve problem instances involving nonlinear nonholonomic dynamics with up to a seven-dimensional state space. The computational savings, evaluated in terms of computation time and storage space, can reach ten orders of magnitude, compared to the standard value iteration algorithm. We have also demonstrated the controller computed by the algorithms on a quadcopter flying quickly through a narrow window, using a motion capture system for state estimation.

Future work will include the extension of the proposed algorithms to more complex problems, including continuous-time partially observable Markov decision processes (POMDPs) and reinforcement learning problems. We will also study the convergence properties of the proposed algorithms in

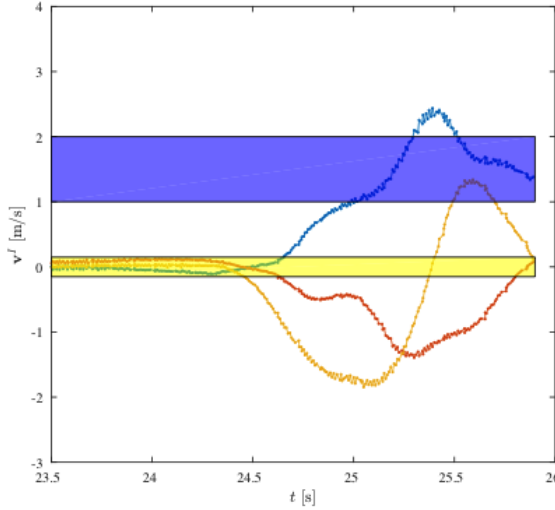
more detail. In particular, we will investigate easily verifiable conditions for which solution ranks can be bounded and therefore algorithm convergence can be proven. If successful, these problems can be shown to be solvable in polynomial time. Finally, we will consider other application domains, such as distributed control systems. We conjecture that the power of the proposed algorithms is greatest when the underlying dynamical system is loosely coupled. We will investigate this conjecture in computational experiments.

### ACKNOWLEDGEMENTS

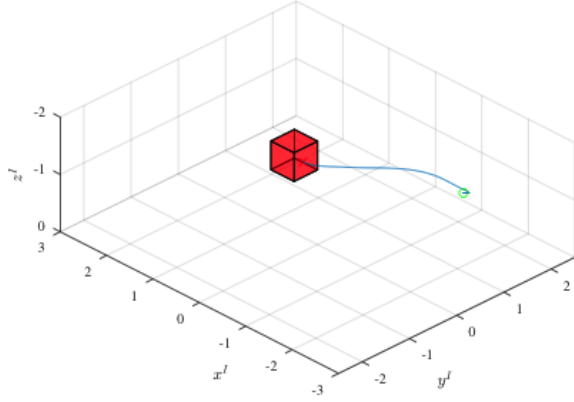
We thank Ezra Tal for pointing out the proof for Lemma 1. We thank Fabian Riether for help in conducting the experiments on an autonomous quadcopter that he has developed in his masters thesis [54].

### REFERENCES

- [1] BATTLES, Z., AND TREFETHEN, L. N. An extension of MATLAB to continuous functions and operators. *SIAM Journal on Scientific Computing* 25, 5 (January 2004), 1743–1770.
- [2] BELLMAN, R. E. *Adaptive Control Processes: A Guided Tour*. Princeton university press, 1961.
- [3] BERTSEKAS, D. P. *Dynamic Programming and Optimal Control*. Athena Scientific Belmont, 2007.
- [4] BERTSEKAS, D. P. *Approximate Dynamic Programming*. Athena Scientific, 2011.
- [5] BERTSEKAS, D. P. *Dynamic Programming and Optimal Control*, 4th ed., vol. II. Athena Scientific, 2012.
- [6] BERTSEKAS, D. P. *Abstract dynamic programming*. Athena Scientific, Belmont, MA, 2013.
- [7] BERTSEKAS, D. P., AND TSITSIKLIS, J. N. *Neuro-Dynamic Programming*, vol. 3 of *Optimization and Neural Computation Series*. Athena Scientific, 1996.
- [8] BIGONI, D., ENGSIG-KARUP, A. P., AND MARZOUK, Y. M. Spectral Tensor-Train decomposition. *SIAM Journal on Scientific Computing* 38, 4 (January 2016), A2405–A2439.



(a) velocities. blue:  $v_x^I$ , red:  $v_y^I$ , yellow:  $v_z^I$ , with corresponding velocity-target regions.



(b) 3D trajectory.

Fig. 20: Motion capture results indicating quadcopter enters into target region [54].

- [9] BRIGGS, W. L., MCCORMICK, S. F., ET AL. *A Multigrid Tutorial*. SIAM, 2000.
- [10] CANNY, J. *The Complexity of Robot Motion Planning*. The MIT Press, 1988.
- [11] CARRILLO, L. R. G., LÓPEZ, A. E. D., LOZANO, R., AND PÉGARD, C. *Quad Rotorcraft Control: Vision-based Hovering and Navigation*. Springer Science & Business Media, 2012.
- [12] CARROLL, J. D., AND CHANG, J.-J. Analysis of individual differences in multidimensional scaling via an n-way generalization of “Eckart-Young” decomposition. *Psychometrika* 35, 3 (September 1970), 283–319.
- [13] CHOW, C.-S., AND TSITSIKLIS, J. An optimal one-way multigrid algorithm for discrete-time stochastic control. *IEEE Transactions on Automatic Control* 36, 8 (1991), 898–914.
- [14] CORY, R., AND TEDRAKE, R. Experiments in fixed-wing UAV perching. In *AIAA Guidance, Navigation and Control Conference and Exhibit* (August 2008), AIAA Reston, VA, American Institute of Aeronautics and Astronautics (AIAA), pp. 1–12.
- [15] DE SILVA, V., AND LIM, L. Tensor rank and the ill-posedness of the best low-rank approximation problem. *SIAM Journal on Matrix Analysis and Applications* 30, 3 (January 2008), 1084–1127.
- [16] DOLGOV, S. V. TT-GMRES: solution to a linear system in the structured tensor format. *Russian Journal of Numerical Analysis and Mathematical Modelling* 28, 2 (January 2013).
- [17] FALLON, M., KUINDERSMA, S., KARUMANCHI, S., ANTONI, M., SCHNEIDER, T., DAI, H., D’ARPINO, C. P., DEITS, R., DICICCO, M., FOURIE, D., KOOLEN, T., MARION, P., POSA, M., VALENZUELA, A., YU, K.-T., SHAH, J., IAGNEMMA, K., TEDRAKE, R., AND TELLER, S. An Architecture for Online Affordance-based Perception and Whole-body Planning. *Journal of Field Robotics* 32, 2 (Oct. 2014), 229–254.
- [18] FENG, S., WHITMAN, E., XINJILEFU, X., AND ATKESON, C. G. Optimization-based Full Body Control for the DARPA Robotics Challenge. *Journal of Field Robotics* 32, 2 (Jan. 2015), 293–312.
- [19] FLEMING, W. H., AND SONER, H. M. *Controlled Markov Processes and Viscosity Solutions*. Springer New York, 2006.
- [20] GOREINOV, S. A., TYRTYSHNIKOV, E. E., AND ZAMARASHKIN, N. L. A theory of pseudoskeleton approximations. *Linear Algebra and its Applications* 261, 1 (1997), 1–21.
- [21] GORODETSKY, A.  $C^3$ : Compressed Continuous Computation library. <https://github.com/goroda/Compressed-Continuous-Computation>.
- [22] GORODETSKY, A. Stochastic optimal control with compressed continuous computation (C3SC). <https://github.com/goroda/c3sc>.
- [23] GORODETSKY, A., KARAMAN, S., AND MARZOUK, Y. Efficient high-dimensional stochastic optimal motion control using Tensor-Train decomposition. In *Robotics: Science and Systems XI* (Rome, Italy, July 2015), Robotics: Science and Systems Foundation.
- [24] GORODETSKY, A., KARAMAN, S., AND MARZOUK, Y. Function-Train: A continuous analogue of the Tensor-Train decomposition. *arXiv preprint arXiv:1510.09088* (2015).
- [25] HACKBUSCH, W. *Tensor Spaces and Numerical Tensor Calculus*. Springer-Verlag, Berlin, 2012.
- [26] HACKBUSCH, W., AND KÜHN, S. A new scheme for the tensor representation. *Journal of Fourier Analysis and Applications* 15, 5 (October 2009), 706–722.
- [27] HASHEMI, B., AND N., T. L. Chebfun in three dimensions, 2016. [http://people.maths.ox.ac.uk/trefethen/Chebfun3\\_submission.pdf](http://people.maths.ox.ac.uk/trefethen/Chebfun3_submission.pdf).
- [28] HÅSTAD, J. Tensor rank is NP-complete. *Journal of Algorithms* 11, 4 (December 1990), 644–654.
- [29] HOROWITZ, M. B., DAMLE, A., AND BURDICK, J. W. Linear Hamilton Jacobi Bellman equations in high dimensions. In *53rd IEEE Conference on Decision and Control* (December 2014), Institute of Electrical & Electronics Engineers (IEEE), pp. 5880–5887.
- [30] HSU, D., LATOMBE, J., AND MOTWANI, R. Path Planning in Expansive Configuration Spaces. In *IEEE Conference on Robotics and Automation* (1997), pp. 2719–2726.
- [31] HSU, D., LATOMBE, J. P., AND KURNIAWATI, H. On the probabilistic foundations of probabilistic roadmap planning. 627–643.
- [32] KARAMAN, S., AND FRAZZOLI, E. Sampling-based Algorithms for Optimal Motion Planning. *International Journal of Robotics Research* 30, 7 (2011), 846–894.
- [33] KAVRAKI, L., KOLOUNTZAKIS, M., AND LATOMBE, J. Analysis of probabilistic roadmaps for path planning. 166–171.
- [34] KAVRAKI, L., SVESTKA, P., LATOMBE, J., AND OVERMARS, M. H. Probabilistic roadmaps for path planning in high-dimensional configuration spaces.
- [35] KOLDA, T. G., AND BADER, B. W. Tensor decompositions and applications. *SIAM Review* 51, 3 (August 2009), 455–500.
- [36] KRUSKAL, J. B. *Multiway data analysis*. Amsterdam, North-Holland, 1989, ch. Rank, decomposition, and uniqueness for 3-way and N-way arrays.
- [37] KUSHNER, H. J. *Probability methods for approximations in stochastic control and for elliptic equations*, vol. 129 of *Mathematics in science and engineering*. New York: Academic Press, 1977.
- [38] KUSHNER, H. J. Numerical methods for stochastic control problems in continuous time. *SIAM Journal on Control and Optimization* 28, 5 (September 1990), 999–1048.
- [39] KUSHNER, H. J., AND DUPUIS, P. G. *Numerical methods for stochastic control problems in continuous time*. Springer, 2001.
- [40] LATOMBE, J. *Robot motion planning*. Springer, 1991.
- [41] LAVALLE, S. M. *Planning algorithms*. Cambridge Univ Pr, May 2006.
- [42] LAVALLE, S. M., AND KUFFNER, J. Randomized kinodynamic planning. 378–400.

- [43] MAHONEY, M. W., AND DRINEAS, P. CUR matrix decompositions for improved data analysis. *Proceedings of the National Academy of Sciences* 106, 3 (January 2009), 697–702.
- [44] MELLINGER, D., MICHAEL, N., AND KUMAR, V. Trajectory Generation and Control for Precise Aggressive Maneuvers with Quadrotors. *International Journal of Robotics Research* 31, 5 (2012), 664–674.
- [45] MOORE, J., AND TEDRAKE, R. Control synthesis and verification for a perching uav using lqr-trees. In *CDC* (2012), Citeseer, pp. 3707–3714.
- [46] OKSENDAL, B. *Stochastic Differential Equations: An Introduction with Applications*. Springer Verlag, 2003.
- [47] OSELEDETS, I. V. Tensor-Train decomposition. *SIAM Journal on Scientific Computing* 33, 5 (January 2011), 2295–2317.
- [48] OSELEDETS, I. V., AND DOLGOV, S. V. Solution of linear systems and matrix inversion in the TT-format. *SIAM Journal on Scientific Computing* 34, 5 (January 2012), A2718–A2739.
- [49] OSELEDETS, I. V., AND TYRTYSHNIKOV, E. TT-cross approximation for multidimensional arrays. *Linear Algebra and its Applications* 432, 1 (January 2010), 70–88.
- [50] PLATTE, R. B., AND TREFETHEN, L. N. Chebfun: A new kind of numerical computing. In *Progress in Industrial Mathematics at ECMI 2008*. Springer Science + Business Media, 2010, pp. 69–87.
- [51] POWELL, W. B. *Approximate Dynamic Programming: Solving the Curses of Dimensionality*, vol. 703. John Wiley & Sons, 2007.
- [52] PRAJNA, S., PARRILO, P. A., AND RANTZER, A. Nonlinear Control Synthesis by Convex Optimization. *IEEE Transactions on Automatic Control* 49, 2 (Feb. 2004), 310–314.
- [53] RATLIFF, N., ZUCKER, M., BAGNELL, J. A., AND SRINIVASA, S. CHOMP: Gradient Optimization Techniques for Efficient Motion Planning. In *IEEE International Conference on Robotics and Automation* (Mar. 2009), pp. 1–6.
- [54] RIETHER, F. Agile quadrotor maneuvering using tensor-decomposition-based globally optimal control and onboard visual-inertial estimation. Masters Thesis, Massachusetts Institute of Technology, 2016.
- [55] ROBERTS, J. W., CORY, R., AND TEDRAKE, R. On the controllability of fixed-wing perching. In *2009 American Control Conference* (2009), IEEE, Institute of Electrical & Electronics Engineers (IEEE), pp. 2018–2023.
- [56] SCIAVICCO, L., AND SICILIANO, B. *Modeling and Control of Robot Manipulators*. Springer, 2000.
- [57] TEDRAKE, R., MANCHESTER, I. R., TOBENKIN, M., AND ROBERTS, J. W. LQR-trees: Feedback Motion Planning via Sums-of-Squares Verification. 1038–1052.
- [58] TOWNSEND, A., AND TREFETHEN, L. N. An extension of Chebfun to two dimensions. *SIAM Journal on Scientific Computing* 35, 6 (January 2013), C495–C518.
- [59] TROTTEMBERG, U., OOSTERLEE, C. W., AND SCHULLER, A. *Multi-grid*. Academic press, 2000.
- [60] TSITSIKLIS, J. Efficient algorithms for globally optimal trajectories. *IEEE Transactions on Automatic Control* 40, 9 (1995), 1528–1538.
- [61] TYRTYSHNIKOV, E. E. Incomplete cross approximation in the mosaic-skeleton method. *Computing* 64, 4 (June 2000), 367–380.
- [62] YANG, I., MORZFELD, M., TOMLIN, C. J., AND CHORIN, A. J. Path integral formulation of stochastic optimal control with generalized costs. *IFAC Proceedings Volumes* 47, 3 (2014), 6994–7000.
- [63] ZUCKER, M., RATLIFF, N., DRAGAN, A. D., PIVTORAIKO, M., KLINGENSMITH, M., DELLIN, C. M., BAGNELL, J. A., AND SRINIVASA, S. S. CHOMP: Covariant Hamiltonian optimization for motion planning. *The International Journal of Robotics Research* 32, 9-10 (Sept. 2013), 1164–1193.